



CORRECTED
EDITION



MIXED-SIGNAL DESIGN SEMINAR

MIXED SIGNAL PROCESSING DESIGN SEMINAR

INTRODUCTION TO MIXED SIGNAL PROCESSING OF
REAL-WORLD SIGNALS AND SIGNAL CONDITIONING

1

LINEAR AND NON-LINEAR ANALOG SIGNAL PROCESSING

2

FUNDAMENTALS OF SAMPLED DATA SYSTEMS

3

ADCs FOR DSP APPLICATIONS

4

DACs FOR DSP APPLICATIONS

5

SIGMA-DELTA ADCs AND DACs

6

DIGITAL SIGNAL PROCESSING TECHNIQUES

7

DSP HARDWARE

8

INTERFACING ADCs AND DACs TO
DIGITAL SIGNAL PROCESSORS

9

MIXED SIGNAL PROCESSING APPLICATIONS

10

MIXED SIGNAL CIRCUIT TECHNIQUES

11

INDEX



ACKNOWLEDGEMENTS

Thanks are due the many technical staff members of Analog Devices in Engineering and Marketing who provided invaluable inputs during the preparation of this seminar. Particular credit is due Bob Fine for his contributions to the DSP material, James Bryant for the Techniques section, and Hank Zumbahlen for the analog filtering section. Principal technical reviewers were James Bryant, Bob Clarke, Bob Fine, Doug Grant, Chris Hyde, Ron Mancini, Paul Brokaw, and Hank Zumbahlen. Personal thanks goes to Dave Kress and Chris Hyde who both provided much-needed guidance and encouragement during this project.

Layout, typesetting, and many of the illustrations were prepared by Linda Grimes Brandon of Brandon's WordService. Camera-ready artwork was supplied by Ernie Lehtonen of the Analog Devices' art department. Printing was done by R. R. Donnelley and Sons, Inc.

Walt Kester
1991

Copyright © 1991 by Analog Devices, Inc.
Printed in the United States of America

All rights reserved. This book, or parts thereof, must not be reproduced in any form without permission of the copyright owner.

Information furnished by Analog Devices, Inc., is believed to be accurate and reliable. However, no responsibility is assumed by Analog Devices, Inc., for its use.

Analog Devices, Inc., makes no representation that the interconnections of its circuits as described herein will not infringe on existing or future patent rights, nor do the descriptions contained herein imply the granting of licenses to make, use, or sell equipment constructed in accordance therewith.

Specifications are subject to change without notice.

ISBN-0-916550-08-7

ANALOG DEVICES MIXED SIGNAL PROCESSING DESIGN SEMINAR

SECTION I

INTRODUCTION TO MIXED SIGNAL PROCESSING OF REAL-WORLD SIGNALS AND SIGNAL CONDITIONING

ORIGINS OF REAL-WORLD SIGNALS AND THEIR UNITS OF MEASUREMENTS

REASONS FOR PROCESSING REAL-WORLD SIGNALS

GENERATION OF REAL-WORLD SIGNALS

**METHODS AND TECHNOLOGIES AVAILABLE FOR PROCESSING REAL-WORLD
SIGNALS**

ANALOG VERSUS DIGITAL SIGNAL PROCESSING

AMPLIFIERS USED AS SIGNAL CONDITIONERS

DIFFERENTIAL AND INSTRUMENTATION AMPLIFIERS

LINE DRIVERS AND RECEIVERS

ISOLATION AMPLIFIERS

SECTION II

LINEAR AND NON-LINEAR ANALOG SIGNAL PROCESSING

AMPLIFIERS USED AS ANALOG SIGNAL PROCESSORS

DISK DRIVE READ AMPLIFIERS

ANALOG MULTIPLIERS

RMS TO DC CONVERTERS

LOGARITHMIC AMPLIFIERS

VARIABLE GAIN AMPLIFIER (ULTRASOUND APPLICATION)

PASSIVE AND ANALOG FILTERING:

ANTIALIASING FILTER DESIGN EXAMPLE, A PROGRAMMABLE STATE VARIABLE FILTER, SEVEN-POLE FDNR 20KHZ ANTIALIASING FILTER, WIDEBAND SALLEN-KEY FILTER

SECTION III

FUNDAMENTALS OF SAMPLED DATA SYSTEMS

A TYPICAL DSP SAMPLED DATA SYSTEM

DISCRETE TIME SAMPLING OF ANALOG SIGNALS

SELECTION OF ANTIALIASING FILTER

OVERSAMPLING AND DECIMATION

UNDERSAMPLING AND ITS APPLICATIONS

EFFECTS OF FINITE AMPLITUDE RESOLUTION DUE TO QUANTIZATION

QUANTIZATION THEORY, SIGNAL TO NOISE RATIO, AND EFFECTIVE BITS

SELECTION OF ADC RESOLUTION BASED ON SIGNAL DYNAMIC RANGE

ADC STATIC TRANSFER CHARACTERISTICS

DAC STATIC TRANSFER CHARACTERISTICS

ADC DYNAMIC PERFORMANCE:

SIGNAL TO NOISE RATIO AND EFFECTIVE BITS, PEAK SPURIOUS, PEAK HARMONIC

CONTENT, SPURIOUS FREE DYNAMIC RANGE, TOTAL HARMONIC DISTORTION, FULL POWER BANDWIDTH, FULL-LINEAR BANDWIDTH, INTERMODULATION DISTORTION (IMD), AC LINEARITY PLOTS USING HISTOGRAMS, APERTURE DELAY TIME (OR EFFECTIVE APERTURE DELAY TIME), APERTURE JITTER, TRANSIENT RESPONSE OR SETTLING TIME, OVERVOLTAGE RECOVERY

DAC DYNAMIC PERFORMANCE:

SETTLING TIME, GLITCH IMPULSE AREA, HARMONIC DISTORTION, DEGLITCHING DACs USING SHAs, $\text{SIN}(X)/X$ FREQUENCY ROLLOFF EFFECT

SWITCHED CAPACITOR FILTERS

SECTION IV

ADCs FOR DSP APPLICATIONS

SUCCESSIVE APPROXIMATION ADCs

FLASH ADCs

SUBRANGING ADCs

INTEGRATING (DUAL SLOPE) ADCs

SECTION V

DACs FOR DSP APPLICATIONS

DAC ARCHITECTURES

GLITCH REDUCTION BY SEGMENTATION

GLITCH REDUCTION BY DIGITAL OFFSET

DEGLITCHING DACs WITH TRACK-AND-HOLDS

MULTIPLYING DACs

LOGDACs

SECTION VI

SIGMA-DELTA ADCs AND DACs

SIGMA-DELTA OVERVIEW

OVERSAMPLING

SIGMA-DELTA MODULATORS AND QUANTIZATION NOISE SHAPING

DIGITAL FILTERING AND DECIMATION

IDLING PATTERNS AND TONAL CONSIDERATIONS FOR SIGMA-DELTA ADCs

HIGHER ORDER MODULATOR LOOPS

DESCRIPTION OF AD1879 18 BIT SIGMA-DELTA AUDIO ADC

SIGMA-DELTA ADCs FOR LOW FREQUENCY MEASUREMENT APPLICATIONS

SIGMA-DELTA DACs

THE ADSP-28msp02 SIGMA-DELTA CODEC

MULTISTAGE NOISE SHAPING (MASH) SIGMA-DELTA CONVERTERS

MULTI-BIT SIGMA-DELTA CONVERTERS

SIGMA-DELTA SUMMARY

SECTION VII

DIGITAL SIGNAL PROCESSING TECHNIQUES

DIGITAL FILTERING:

FINITE IMPULSE RESPONSE (FIR) FILTERS, THE DUALITY OF THE TIME AND FREQUENCY DOMAIN, FIR FILTER IMPLEMENTATION IN DSP HARDWARE USING CIRCULAR BUFFERING, FIR FILTER DESIGN TECHNIQUES, FILTER DESIGN USING CAD TECHNIQUES, DESIGN EXAMPLE FOR AN FIR DIGITAL AUDIO FILTER USING CAD PROGRAM, INSURING LINEAR PHASE IN FIR FILTERS, DECIMATION USING FIR FILTERS, INFINITE IMPULSE RESPONSE (IIR) DIGITAL FILTERS, SUMMARY: FIR VERSUS IIR FILTERS

FAST FOURIER TRANSFORMS:

FFT HARDWARE IMPLEMENTATION, FFT DESIGN CONSIDERATIONS, SPECTRAL LEAKAGE AND WINDOWING, DATA SCALING AND BLOCK FLOATING POINT

SECTION VIII

DSP HARDWARE

RISC VERSUS CISC VERSUS DSP ARCHITECTURES

RISC AND DSP APPLICATIONS

DSP PROCESSOR REQUIREMENTS:

FAST ARITHMETIC, ZERO OVERHEAD LOOPING, EXTENDED DYNAMIC RANGE, DUAL OPERAND FETCH, CIRCULAR BUFFERS

ADSP-2101 MICROCOMPUTER GENERAL DESCRIPTION

ADSP-2101 ARCHITECTURE OVERVIEW:

ARITHMETIC LOGIC UNIT (ALU), MULTIPLIER/ACCUMULATOR (MAC),
SHIFTER, DATA ADDRESS GENERATORS (DAGS), PROGRAM SEQUENCER,
SERIAL PORTS, SYSTEM INTERFACE

DEVELOPMENT SYSTEM

SECTION IX

INTERFACING ADCs AND DACs TO DIGITAL SIGNAL PROCESSORS

**PARALLEL INTERFACING TO DSP PROCESSORS: READING DATA FROM
MEMORY-MAPPED PERIPHERAL ADCs**

PARALLEL ADC TO DSP INTERFACE

**PARALLEL INTERFACING TO DSP PROCESSORS: WRITING DATA TO
MEMORY-MAPPED DACs**

PARALLEL DAC TO DSP INTERFACE

SERIAL INTERFACING TO DSP PROCESSORS:

SERIAL ADC TO DSP INTERFACE, SERIAL DAC TO DSP INTERFACE

INTERFACING I/O PORTS AND CODECS TO DSPs

SERIAL VERSUS PARALLEL DSP INTERFACE SUMMARY

SECTION X

MIXED SIGNAL PROCESSING APPLICATIONS

HIGH PERFORMANCE MODEMS:

V.32 MODEM OVERVIEW, V.32 MODEM TRANSMITTER, V.32 MODEM
RECEIVER, I/O PORTS AND CODECS FOR V.32 MODEMS

DIGITAL MOBILE RADIO:

OVERVIEW, THE GSM SYSTEM, SPEECH CODEC, DISCONTINUOUS
TRANSMISSION (DTX), GSM SYSTEM UPCONVERSION AND
DOWNCONVERSION

DIGITAL AUDIO STUDIO RECORDING

COMPACT DISC (CD) PLAYER ELECTRONICS

SECTION XI MIXED SIGNAL CIRCUIT TECHNIQUES

INTRODUCTION

RESISTANCE:

RESISTANCE OF CONDUCTORS, SKIN EFFECT, VOLTAGE DROP IN SIGNAL LEADS-KELVIN FEEDBACK, LEAKAGE IN INSULATORS, GUARD RINGS, ELECTROSTATIC DAMAGE, PARASITIC EFFECTS IN RESISTORS, INDUCTANCE, THERMOELECTRIC EFFECTS, STABILITY AND MATCHING, VOLTAGE VARIATION OF RESISTANCE, JOHNSON NOISE

CAPACITANCE:

STRAY CAPACITANCE, FARADAY SHIELDS, NOISE, PARASITIC EFFECTS IN CAPACITORS, CAPACITOR LEAKAGE, SERIES/LOSS RESISTANCE, INDUCTANCE OF CAPACITORS, DIELECTRIC ABSORPTION

INDUCTANCE:

STRAY INDUCTANCE, MUTUAL INDUCTANCE, RINGING, PARASITIC EFFECTS IN INDUCTORS, QUALITY FACTOR (Q)

GROUNDING AND SIGNAL ROUTING:

SIGNAL RETURN CURRENTS, GROUND NOISE AND GROUND LOOPS, STAR (MECCA) GROUNDS, SEPARATE ANALOG AND DIGITAL GROUNDS, GROUND PLANES, TRANSMISSION LINES, SYSTEM GROUNDS, SIGNAL ROUTING

POWER SUPPLIES:

POWER SUPPLY NOISE, SWITCHING-MODE POWER SUPPLIES

ELECTROMAGNETIC INTERFERENCE:

RADIO FREQUENCY INTERFERENCE, PHOTOELECTRIC EFFECTS

LOGIC:

FAN-OUT, TIMING VARIATIONS, SAMPLING CLOCK NOISE, LOGIC NOISE

PROBLEM AREAS:

LIMITATIONS OF SPICE MODELLING, SOCKETS, PROTOTYPING HIGH PERFORMANCE ANALOG CIRCUITRY

INDEX

MIXED SIGNAL PROCESSING DESIGN SEMINAR

INTRODUCTION TO MIXED SIGNAL PROCESSING OF
REAL-WORLD SIGNALS AND SIGNAL CONDITIONING

1

LINEAR AND NON-LINEAR ANALOG SIGNAL PROCESSING

2

FUNDAMENTALS OF SAMPLED DATA SYSTEMS

3

ADCs FOR DSP APPLICATIONS

4

DACs FOR DSP APPLICATIONS

5

SIGMA-DELTA ADCs AND DACs

6

DIGITAL SIGNAL PROCESSING TECHNIQUES

7

DSP HARDWARE

8

INTERFACING ADCs AND DACs TO
DIGITAL SIGNAL PROCESSORS

9

MIXED SIGNAL PROCESSING APPLICATIONS

10

MIXED SIGNAL CIRCUIT TECHNIQUES

11

INDEX



SECTION I

INTRODUCTION TO MIXED SIGNAL PROCESSING OF REAL-WORLD SIGNALS AND SIGNAL CONDITIONING

INTRODUCTION TO MIXED SIGNAL PROCESSING OF REAL-WORLD SIGNALS AND SIGNAL CONDITIONING

- **ORIGINS OF REAL-WORLD SIGNALS AND THEIR UNITS OF MEASUREMENT**
- **REASONS FOR PROCESSING REAL-WORLD SIGNALS**
- **GENERATION OF REAL-WORLD SIGNALS**
- **METHODS AND TECHNOLOGIES AVAILABLE FOR PROCESSING REAL-WORLD SIGNALS**
- **ANALOG VERSUS DIGITAL SIGNAL PROCESSING**
- **AMPLIFIERS USED AS SIGNAL CONDITIONERS**
- **DIFFERENTIAL AND INSTRUMENTATION AMPLIFIERS**
- **LINE DRIVERS AND RECEIVERS**
- **ISOLATION AMPLIFIERS**

SECTION I

INTRODUCTION TO MIXED SIGNAL PROCESSING OF REAL-WORLD SIGNALS AND SIGNAL CONDITIONING

ORIGINS OF REAL-WORLD SIGNALS AND THEIR UNITS OF MEASUREMENT

In this seminar, we will primarily be dealing with the processing of *real-world* signals using both analog and digital techniques. Before starting, however, let's look at a few key concepts and definitions required to lay the groundwork for things to come.

Webster's *New Collegiate Dictionary* defines a *signal* as "A detectable (or measur-

able) physical quantity or impulse (as voltage, current, or magnetic field strength) by which messages or information can be transmitted." Key to this definition are the words: *detectable, physical quantity, and information.*

FUNDAMENTAL CHARACTERISTICS OF SIGNALS

- Are Physical Quantities
- Are Measurable
- Contain Information
- All Signals Are Analog
- Some Signals Are Responses to Known Signals Which Act as a Stimulus (i.e. Radar and Ultrasound Return Signals)

Figure 1.1

By their very nature, signals are analog, whether dc, ac, digital levels, or pulses. It is customary, however, to differentiate between *analog* and *digital* signals in the following manner: Analog (or real-world) variables in nature include all measurable physical quantities. In this seminar, *analog* signals are generally limited to electrical variables, their rates of change, and their associated energy or power levels. Sensors and transducers are used to convert other physi-

cal quantities (temperature, pressure, etc.) to electrical signals and vice versa. The entire subject of signal conditioning deals with preparing real-world signals for processing and includes such topics as sensors (temperature and pressure, for example), isolation and instrumentation amplifiers, etc.

Some signals result in response to other signals. A good example is a radar or ultrasound imaging return signal, both of which result from a known transmitted signal.

UNITS OF MEASUREMENT

- Temperature: °C
- Pressure: Newtons/m²
- Mass: kg
- Voltage: Volts
- Current: Amps
- Power: Watts

Figure 1.2

On the other hand, there is another classification of signals, called *digital*, where the actual signal has been conditioned and formatted into a digit. These digital signals may or may not be related to real-world analog variables. In the specific case of Digital Signal Processing (DSP), the analog signal is converted into binary form by a device known as an analog-to-digital converter (ADC). The output of the ADC is a binary representation of the analog signal

and is manipulated arithmetically by the Digital Signal Processor. After processing, the information obtained from the signal may be converted back into analog form using a digital-to-analog converter (DAC).

Another key concept embodied in the definition of *signal* is that there is some kind of *information* contained in the signal. This leads us to the key reason for processing real-world analog signals: the *extraction of information*.

REASONS FOR PROCESSING REAL-WORLD SIGNALS

The primary reason for processing real-world signals is to extract information from them. This information normally exists in the form of signal amplitude (absolute or relative), frequency or spectral content, or timing relationships with respect to other signals. Once the desired information is extracted from the signal, it may be used in a number of ways.

In some cases, it may be desirable to reformat the information contained in a signal. This would be the case in the transmission of a voice signal over a frequency division multiplexed (FDM) telephone system. In this case, analog techniques are

used to “stack” voice channels in the frequency spectrum for transmission via microwave relay or coaxial cable. In the case of a digital transmission link, the analog voice information is first converted into digital using an ADC. The digital information representing the individual voice channels is multiplexed in time (time division multiplexed, or TDM) and transmitted over a serial digital transmission link (as in the T-Carrier system).

Another requirement for signal processing is to *compress* the frequency content of the signal (without losing significant information) then format and transmit the informa-

tion at lower data rates, thereby achieving large reductions in required channel bandwidths. High speed modems and adaptive pulse code modulation systems (ADPCM) make extensive use of data reduction algorithms, as do digital mobile radio systems and High Definition Television (HDTV).

Industrial data acquisition and control systems make use of information extracted from sensors to develop appropriate feedback signals which in turn control the process itself. A block diagram of such a system is

shown in Figure 1.3. Note that these systems require both ADCs and DACs as well as sensors, signal conditioners, and the DSP.

In some cases, the signal containing the information is buried in noise, and the primary objective of signal processing is recovery. Techniques such as filtering, autocorrelation, convolution, etc. are often used to accomplish this task in both the analog and digital domains.

TYPICAL DATA ACQUISITION AND PROCESS CONTROL SYSTEM

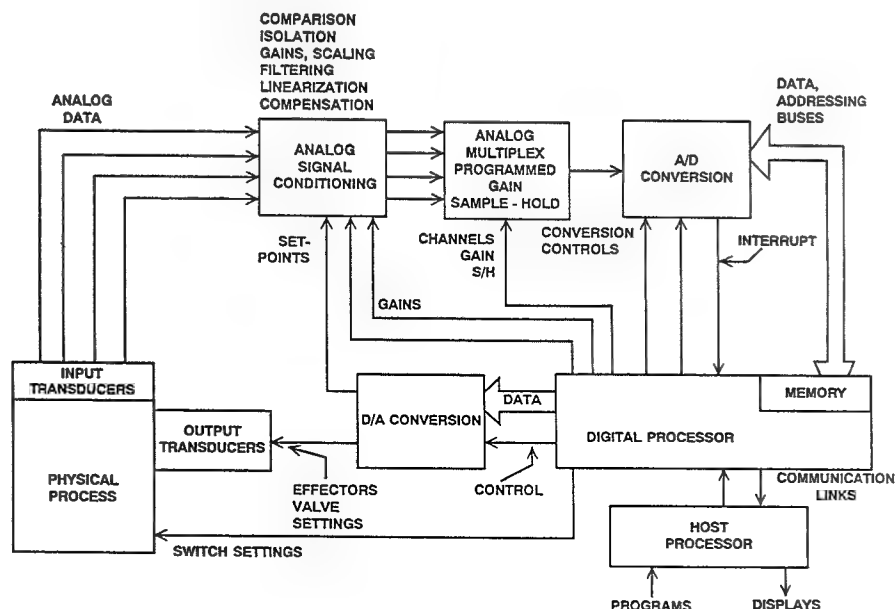


Figure 1.3

REASONS FOR PROCESSING REAL-WORLD SIGNALS

- Extract Information (Amplitude, Frequency, Spectral Content, Timing Relationships)
- Reformat the Signal (FDM, TDM Systems)
- Compression of Data (Modems, Digital Mobile Radio, ADPCM, HDTV)
- Generate Feedback Control Signal (Industrial Process Control)
- Extract Signal from Noise (Filtering, Autocorrelation, Convolution)
- Store Signal Data in Digital Format for Recovery and/or Analysis using DSP Techniques (FFT Analysis)

Figure 1.4

GENERATION OF REAL-WORLD SIGNALS

In most of the above examples (the ones requiring DSP techniques), a DAC is required in order to generate an appropriate analog signal after the DSP has completed its processing function on the converted analog signal. In some cases, however, real world analog signals may be generated directly using DSP and DACs, while omitting the requirement for the front-end ADC. A video example of this is the digital generation of signals which drive videodacs and RAMDACs in raster scan display systems. Artificially synthesized music and speech are

good low frequency examples. In reality, however, the real-world analog signals generated using purely digital techniques do rely on information previously derived from analog signals. In display systems, the data from the display must convey the appropriate information to the operator. In synthesized audio systems, the statistical properties of the sounds being generated have been previously derived using extensive DSP analysis (i.e., sound source, microphone, preamp, ADC, etc.).

DIGITALLY SYNTHESIZED REAL-WORLD SIGNALS

- Graphics Display Systems
- Artificial Synthesis of Sound (Music, Speech)
- These only Require DSP and DACs

Figure 1.5

METHODS AND TECHNOLOGIES AVAILABLE FOR PROCESSING REAL-WORLD SIGNALS

Signals may be processed using analog techniques (analog signal processing, or ASP), digital techniques (digital signal processing, or DSP), or a combination of analog and digital techniques (mixed signal

processing, or MSP). In some cases, the choice of techniques is clear, in others, there is no clear cut choice, and second-order considerations may be used to make the final decision.

PROCESSING REAL-WORLD SIGNALS

- Analog Signal Processing (ASP): Filtering, Amplification, Modulation, Demodulation, Multiplication, Division, Measurement
- Digital Signal Processing (DSP): Filtering, Amplitude Scaling, Modulation, Multiplication
- Mixed Signal Processing (MSP): Analog and Digital Signal Processing Combined in the Same Function: PC Board, Hybrid, or IC, with Implicit *Real-Time* Operation

Figure 1.6

DIGITAL AUDIO STUDIO SYSTEM

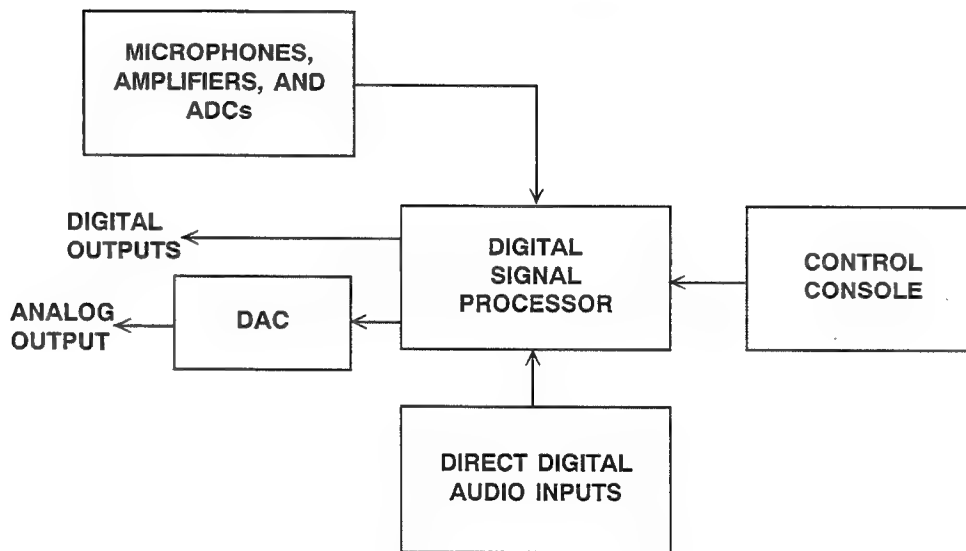


Figure 1.7

With respect to DSP, the factor that distinguishes it from traditional computer analysis of data is its speed and its ability to perform sophisticated digital processing functions. In order to understand the significance of real-time DSP, consider the much simplified digital audio system shown in Figure 1.7. After conversion, all audio processing such as mixing, equalization, filtering, dynamic range control, etc., is handled by the DSP. After processing, the signal is converted back into analog format using a DAC. The ADC sampling rate for such a system is typically 44.1 or 48kSPS. The DSP processing steps, such as digital filtering, must be completed within one cycle of the sampling clock ($23\mu\text{s}$ at 44.1kSPS) in order to “keep up” with the analog signal. This is what is typically

meant by *real-time* DSP operation.

The term *mixed signal processing* implies that *both* analog and digital processing is done as part of the same functional block. This functional block may be implemented in the form of a system, a printed circuit board or hybrid microcircuit, or even in the form of a single integrated circuit chip. In the context of this broad definition, ADCs and DACs are considered to be mixed signal processors, since both analog and digital functions are implemented in each. Recent advances in Very Large Scale Integration (VLSI) technology allow complex digital processing as well as analog processing to be performed on the same chip. The very nature of DSP itself implies that these functions can be performed in *real-time*.

ANALOG VERSUS DIGITAL SIGNAL PROCESSING

Today's engineer faces a challenge in selecting the proper mix of analog and digital techniques to solve the signal processing task at hand. It is impossible to process real-world analog signals using purely digital techniques, since all transducers (thermocouples, strain gages, microphones, piezoelec-

tric crystals, disk drive heads, etc.) are inherently analog elements. Therefore, some sort of signal conditioning circuitry is required in order to prepare the transducer output for further signal processing, whether it be analog or digital. Signal conditioning circuits are, in reality, analog signal proces-

MIXED SIGNAL PROCESSING DESIGN SEMINAR

sors, performing such functions as multiplication (gain), isolation (instrumentation amplifiers and isolation amplifiers), detection in the presence of noise (high common-mode instrumentation amplifiers, line drivers, and

line receivers), dynamic range compression (log amps, LOGDACs, and programmable gain amplifiers), and filtering (both passive and active).

ANALOG SIGNAL CONDITIONING AND PROCESSING

- Amplification (Gain)
- Impedance Transformation
- Removing Common Mode Noise
- Isolation
- Cable Driving and Receiving
- Multiplication of Signals
- Dynamic Range Compression
- Programmable Amplification
- Filtering (Passive and Active)

Figure 1.8

Several methods of accomplishing signal processing are shown in Figure 1.9. The top portion of the figure shows the purely analog approach. The latter two parts of the figure show the DSP approach. Note that once the decision has been made to use DSP techniques, the next decision must be where to place the ADC in the signal path. In general,

as the ADC is moved closer to the actual transducer, more of the analog signal conditioning burden is placed on the ADC. This added ADC complexity may take the form of increased sampling rate, wider dynamic range, higher resolution, input noise rejection, input filtering, etc., all of which imply greater ADC costs. In fact, the probability of

ANALOG AND DIGITAL SIGNAL PROCESSING OPTIONS

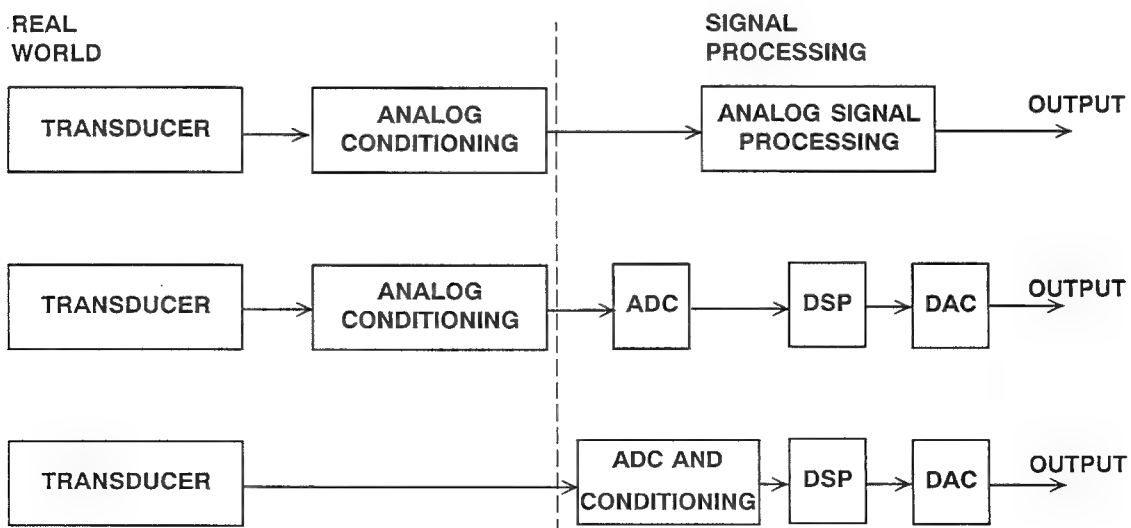


Figure 1.9

finding an ADC that is directly matched to the transducer output and has all the other desired system characteristics is indeed quite remote, except in very special cases such as the simple bimetallic thermal switch.

The system designer should face the fact that some type of signal conditioning will inevitably be required before the actual ADC. Although there are no hard and fast rules available as to where in the signal path to

place the ADC, the current state-of-the-art in sampling ADCs shown in Figure 1.10 may be useful in determining the initial boundary. Usually, the performance requirements on the ADC (and hence the cost) can be relaxed at the expense of additional analog signal conditioning and processing. For instance, a programmable gain amplifier, or a logarithmic amplifier placed ahead of the ADC may reduce the ADC dynamic range requirements.

STATE OF THE ART IN ADCs

Resolution	Sampling Rate
22 bits	1 kSPS
20 bits	4 kSPS
18 bits	50 kSPS
16 bits	500 kSPS
14 bits	10 MSPS
12 bits	25 MSPS
10 bits	75 MSPS
8 bits	500 MSPS

Figure 1.10

In order to understand how to most effectively utilize either analog and/or digital signal processing, the system designer must first understand the capabilities of each. The

following sections will investigate analog elements both as conditioners and as signal processors themselves.

AMPLIFIERS USED AS SIGNAL CONDITIONERS

Operational amplifiers are extremely useful devices for coupling transducer outputs to the signal processor inputs (see Figure 1.11). Op amps can provide gain, impedance transformation, filtering, and level shifting. The amplifier's input characteristics must match the output characteristics of the transducer with respect to impedance, signal level, dynamic range, bandwidth, etc. The output of the amplifier must also match the input characteristics of the signal processor with respect to the same characteristics. If the signal processing chain begins with an ADC, it is extremely important that the signal applied to the ADC

adequately fills the converter's input range without overdrive. Small signals will not fully utilize the dynamic range of the ADC, while signals which are too large will overdrive the converter and cause hard-limiting. In addition, care must be taken to insure that the amplifier does not degrade the performance of the signal processor, especially with respect to dynamic performance specifications such as total harmonic distortion (THD), signal-to-noise ratio, etc. Fortunately, a wide variety of precision dc coupled high speed IC op amps are available to fit almost every application.

THE AMPLIFIER AS A SIGNAL CONDITIONER

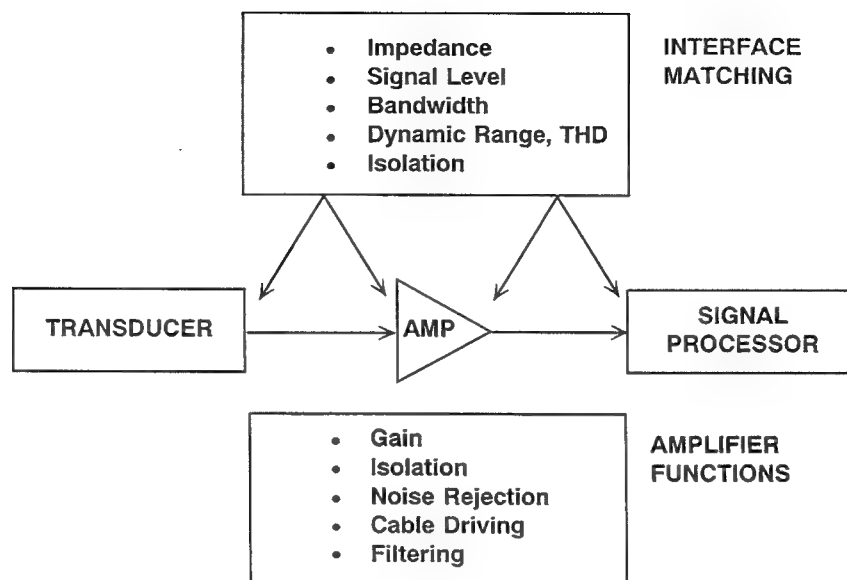


Figure 1.11

KEY OP AMP SPECS

- | | |
|----------------------------|------------------------------|
| ■ Open Loop Gain | ■ Bandwidth |
| ■ Input Impedance | ■ Settling Time |
| ■ Input Offset Voltage | ■ Distortion (THD) |
| ■ Input Bias Currents | ■ Noise: Voltage and Current |
| ■ Output Voltage/Current | ■ Slewrate |
| ■ Temperature Coefficients | |
| ■ Long Term Drift | |

Figure 1.12

Without going into op amp theory in much detail, we should, however, examine a few basic op amp configurations which are suitable to a wide variety of applications. Figure 1.13 shows the two most fundamental op amp configurations: inverting and non-inverting. Both configurations have their

relative strengths and weaknesses. The inverting mode is more commonly used when level shifting is required, but presents an input impedance of R_1 to the input. The non-inverting mode has a high input impedance, but is more awkward to use if level shifting is required (see Figure 1.14).

BASIC OP AMP CONFIGURATIONS

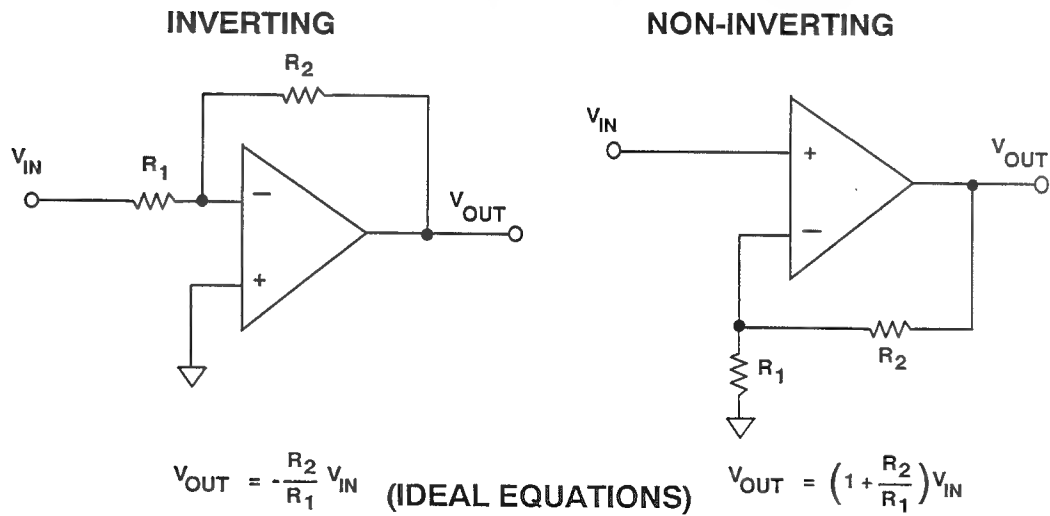


Figure 1.13

OP AMPS AS LEVEL SHIFTERS

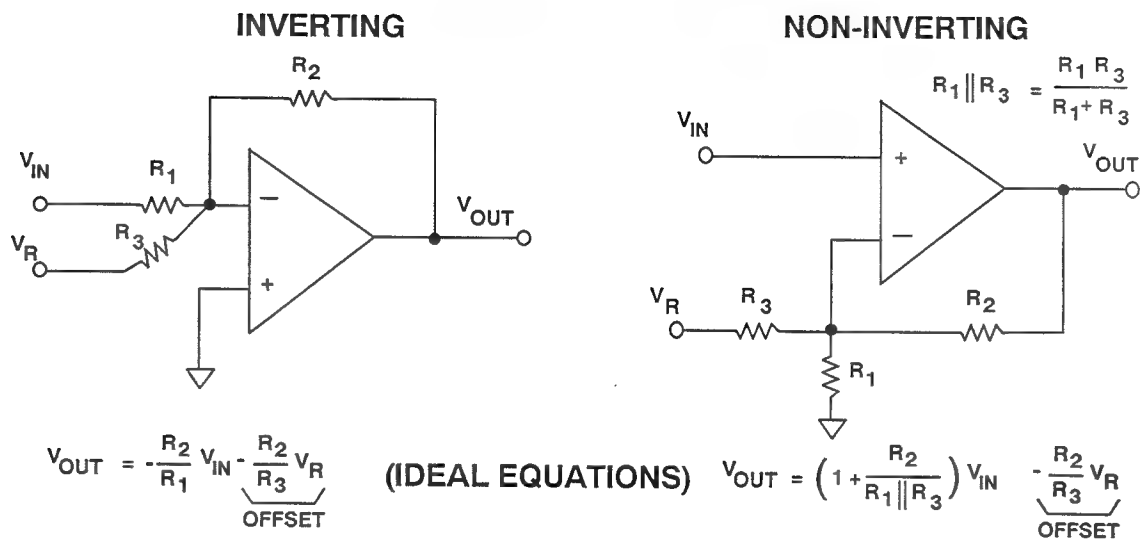


Figure 1.14

DIFFERENTIAL AND INSTRUMENTATION AMPLIFIERS

In many cases, the signal from the transducer is corrupted by the presence of an unwanted common-mode signal such as noise, or ac power line coupling. In these cases, a differential input configuration such as the one shown in Figure 1.15 may be desirable. Signals which are common to both V_1 and V_2 are rejected from the output by the common-mode rejection ratio (CMRR) of the op amp. Unfortunately, CMRR decreases as a function of frequency, and high frequency noise is not rejected. In addition, the CMRR of the differential input amplifier depends upon accurate ratio matching of the four resistors. A mismatch of 0.1% in any of the

four resistors will produce a CMRR of approximately 66dB. Another problem with the simple differential circuit is that V_1 and V_2 see different input impedances, i.e., the input is unbalanced. Even with these sources of error, the differential configuration is quite useful as a line receiver for properly terminated cables, such as 600Ω audio. In this case, the unbalanced input is not a real problem, since the value of R_1 is chosen to be much higher (typically greater than 20kΩ) than the 600Ω line impedance. A discussion of audio line drivers and receivers will follow shortly.

DIFFERENTIAL RECEIVER INPUT CONFIGURATION




Figure 1.15

For true balanced, high impedance inputs, three op amps may be connected to form an *instrumentation amplifier* as shown in Figure 1.16. In this configuration, both V_1 and V_2 see high input impedances, and the input is balanced. The gain of the amplifier is usu-

ally set by an external resistor, R_g , as in the case of the AD521 and AD625. Such circuits are generally used for dedicated fixed-gain applications. Common mode rejection is very high as shown in Figure 1.17. Pin-programmable instrumentation amps, such as the

3 OP-AMP INSTRUMENTATION AMP

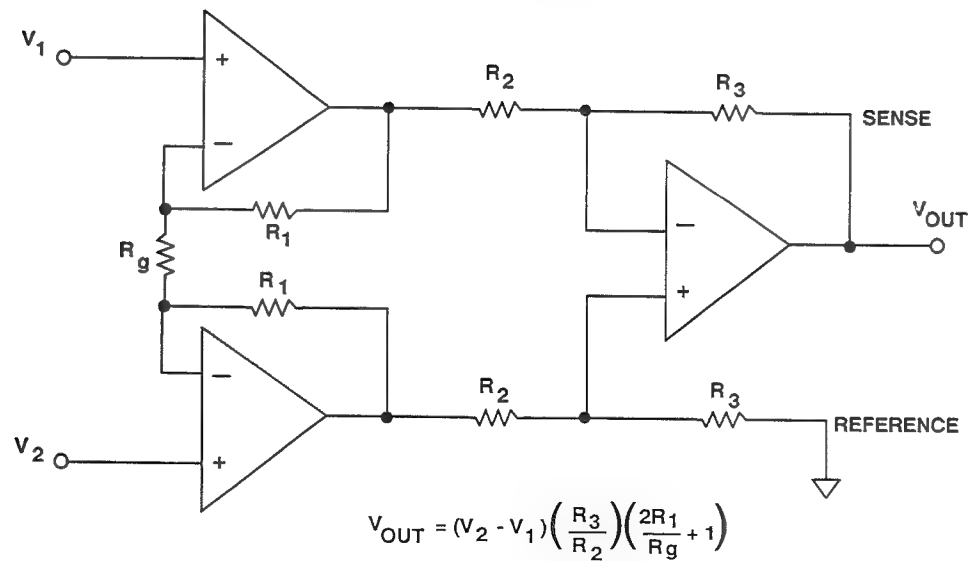


Figure 1.16

INSTRUMENTATION AMPLIFIER USED TO REJECT COMMON MODE VOLTAGE

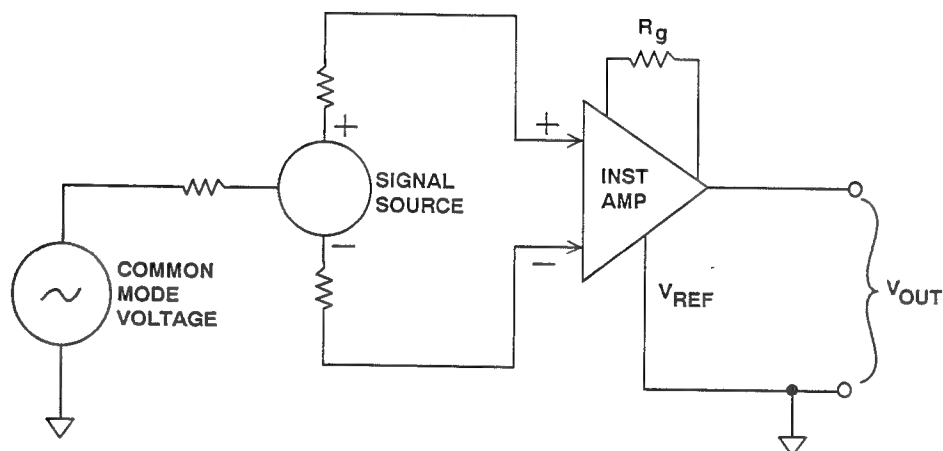


Figure 1.17

HIGH-PERFORMANCE, HIGH-SPEED INSTRUMENTATION AMPLIFIER

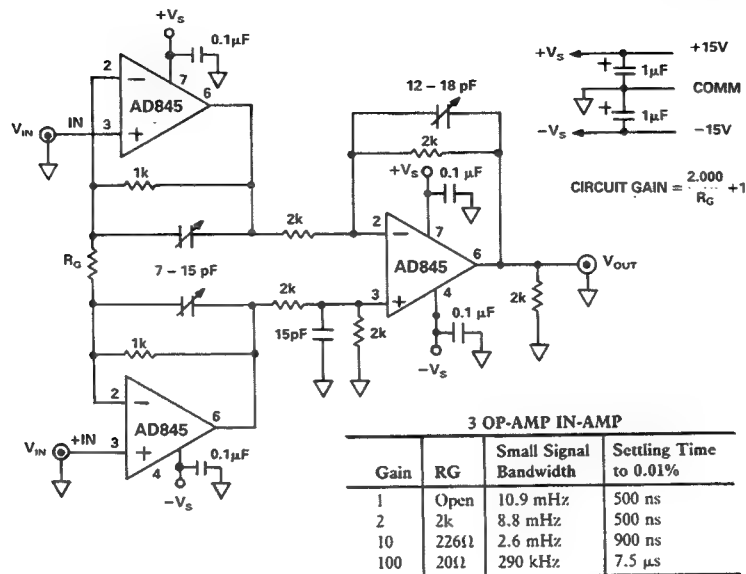



Figure 1.18

AD524 and AD624, have a set of internal resistors; a limited set of fixed gains in the range of 1 to 1000 are chosen by appropriately interconnecting the resistors via external pins. Digitally (or *software*-) programmable instrumentation amps are completely self-contained, with gains set by a 2, 3, or 4 bit digital code as in the case of the AD365 (gains of 1, 10, 100, 500).


A wideband instrumentation amplifier using three AD845 FET input op amps is

shown in Figure 1.18. Gain can be adjusted from 1 to 1000. Low input bias currents and fast settling times are achieved with the AD845. The AD843 FET input op amp may also be used for even higher bandwidth. Note that the bandwidth for the circuit is 10.9MHz at a gain of 1 and 2.6MHz at a gain of 10. Settling time for the entire circuit is 500ns to 0.01% for a 10V step at a gain of 1 as shown in Figure 1.19.

HIGH SPEED INSTRUMENTATION AMP PULSE RESPONSE AND SETTLING TIME



The Pulse Response of the Three Op Amp
Instrumentation Amplifier: Gain = 1,
Horizontal Scale: 0.5µs/Div, Vertical Scale: 5V/Div



Settling Time of the Three Op Amp
Instrumentation Amplifier: Horizontal Scale:
200 ns/Div, Vertical Scale, Pulse Input: 5V/Div;
Output Settling: 1mV/Div

Figure 1.19

LINE DRIVERS AND RECEIVERS

To avoid noise pickup through the cabling, it is desirable to place the signal processor as close to the conditioning circuits as possible. If this is not possible, then the analog signal must be transmitted over cable, either coaxial or shielded twisted pair so that any noise picked up will be common-mode. This requires a balanced line driver at the sending end and a balanced line receiver at the receiving end. A good example of a balanced line driver for audio bandwidths is the Analog Devices/PMI SSM-2142. This device converts a single-ended input signal to a fully balanced, high drive, high output signal pair (see Figure 1.20). The SSM-2142 mimics the fully balanced performance of transformer-based solutions for line driving. However, the SSM-2142 maintains lower distortion

and occupies much less board space than transformers, while achieving comparable common-mode rejection. Since the output stages have balanced impedances due to active laser trimming, hum and noise are rejected over the full audio bandwidth. It is suggested that a suitable differential input amplifier such as the SSM-2141 is used at the receiving end to maintain overall system performance. A functional block diagram of the device along with key specifications are given in Figure 1.21. A typical system application for the driver and receiver combination is shown in Figure 1.22. Typical THD performance at the SSM-2141 single-ended output for a 10V rms output is 0.004% at 1kHz using 500 feet of Belden 8451 cable.

SSM-2142 BALANCED AUDIO LINE DRIVER

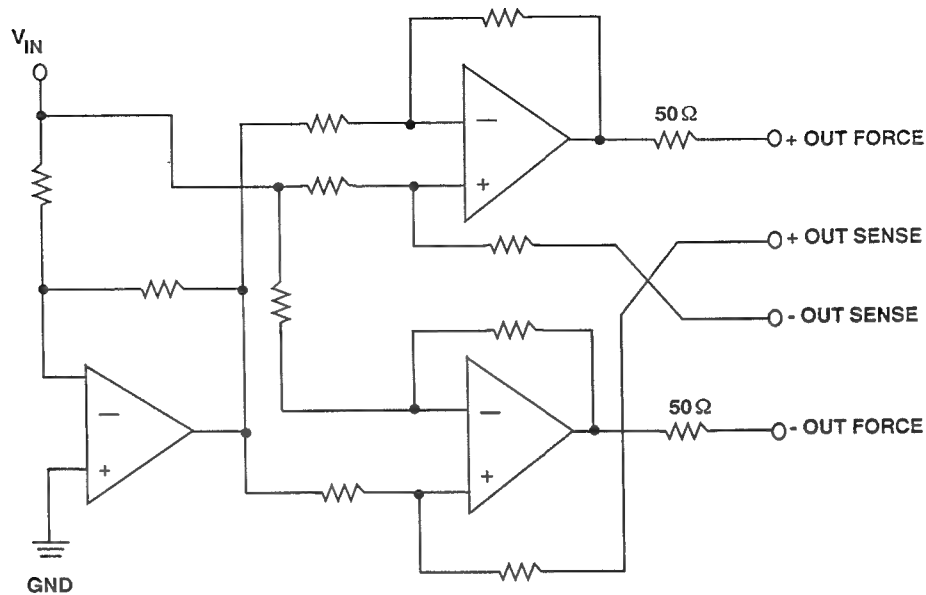



Figure 1.20

SM-2141 HIGH COMMON-MODE REJECTION DIFFERENTIAL AUDIO LINE RECEIVER




■ CMRR: 100dB @ 60 Hz
70dB @ 20 kHz
62dB @ 40 kHz

■ THD: 0.001%

■ BANDWIDTH: 3MHz

Figure 1.21

AUDIO LINE DRIVER AND RECEIVER (SSM-2142, SSM-2141)



■ THD: 0.004% FOR $V_{OUT} = 10V_{rms}$

Figure 1.22

Differential techniques such as those described above become difficult if not impossible at video frequencies primarily because of the lack of suitable line drivers and receivers having good common mode rejection and low distortion at high frequencies. Furthermore, video signals must be transmitted over properly terminated cable in order to avoid standing waves and distortion. Typical configurations for high speed cable driving are shown in Figure 1.23. Note that the cables are terminated both at the source and at the load to minimize reflections. This implies that a gain of at least two is required

in order to restore the signal to its original level at the receiving end of the cable. Figure 1.24 shows a video line driver circuit using the AD829 which has been optimized for bandwidth flatness and low differential gain and phase. This circuit will drive reverse-terminated 75 Ω video cable to standard video levels (1V p-p) with 0.1dB gain flatness to 30MHz with only 0.02° and 0.02% differential phase and gain at the 4.43MHz PAL color subcarrier frequency. This level of performance meets the requirements for high-definition video displays and test equipment.

NONINVERTING AND INVERTING CABLE DRIVERS

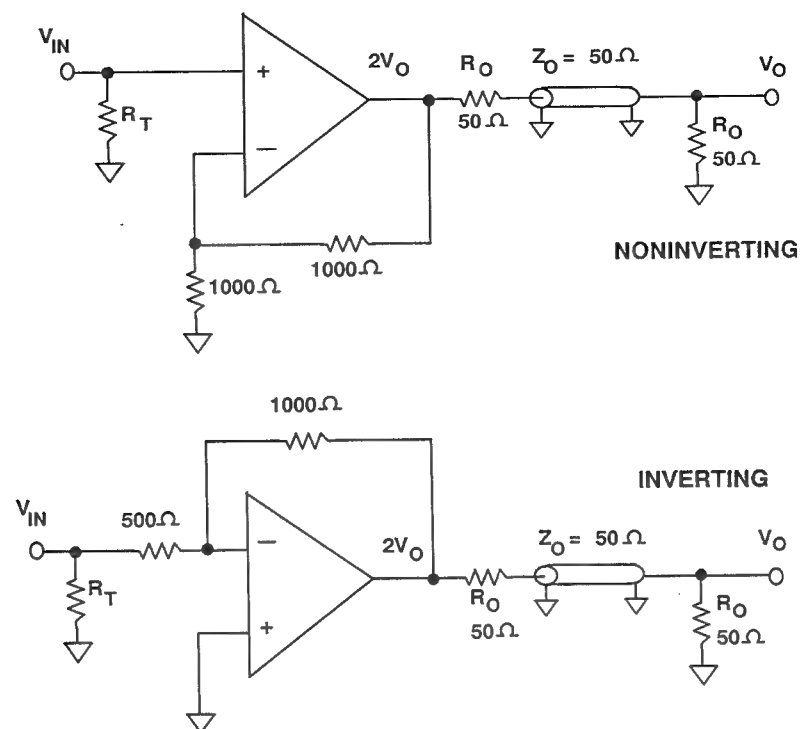


Figure 1.23

LOW DIFFERENTIAL GAIN AND PHASE VIDEO LINE DRIVER: AD829

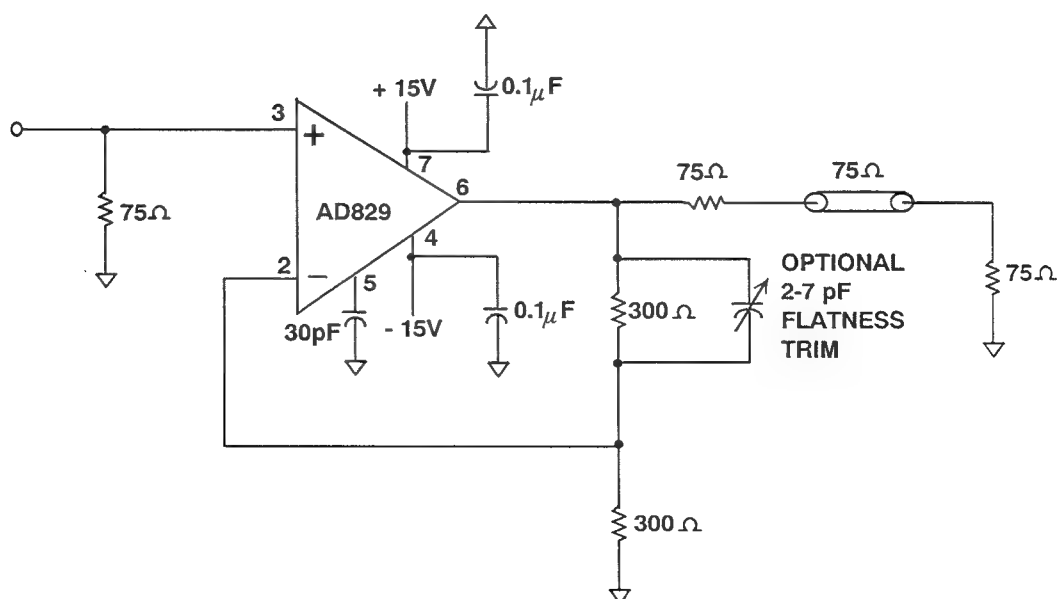


Figure 1.24

ISOLATION AMPLIFIERS

The isolation amplifier has an input circuit that is galvanically isolated from the power supply and the output circuit. In addition, there is minimal capacitance between the input and the rest of the device. Therefore, there is no possibility for dc current flow, and minimum ac coupling. Isolation amplifiers are intended for applications requiring safe, accurate measurement of low frequency voltage or current (up to about 100kHz) in the presence of high common-mode voltage (to thousands of volts) with high common mode rejection; line-receiving of signals transmitted at high impedance in noisy environments; and for safety in general-purpose measurements where dc and line-frequency leakage must be maintained at levels well below certain mandated minima. Principle applications are in electrical environments of the kind associated with medical equipment, conventional and nuclear power plants, automatic test equipment, and industrial process control systems.

In the basic two-port form, the output and power circuits are not isolated from one

another. In the three-port isolator shown in Figure 1.25, the input circuits, output circuits, and power source are all isolated from one another. The figure shows the circuit architecture of a self-contained isolator, the AD210. An isolator of this type requires power from a two-terminal dc power supply. An internal oscillator (50kHz) converts the dc power to ac, which is transformer-coupled to the shielded input section, then converted to dc for the input stage and the auxiliary power output. The ac carrier is also modulated by the amplifier output, transformer-coupled to the output stage, demodulated by a phase-sensitive demodulator (using the carrier as the reference), filtered and buffered, using isolated dc power derived from the carrier. The AD210 allows the user to select gains from 1 to 100 using an external resistor. Bandwidth is 20kHz, and voltage isolation is 2500V rms (continuous) and $\pm 3500\text{V}$ peak (continuous).

AD210 THREE-PORT ISOLATION AMPLIFIER

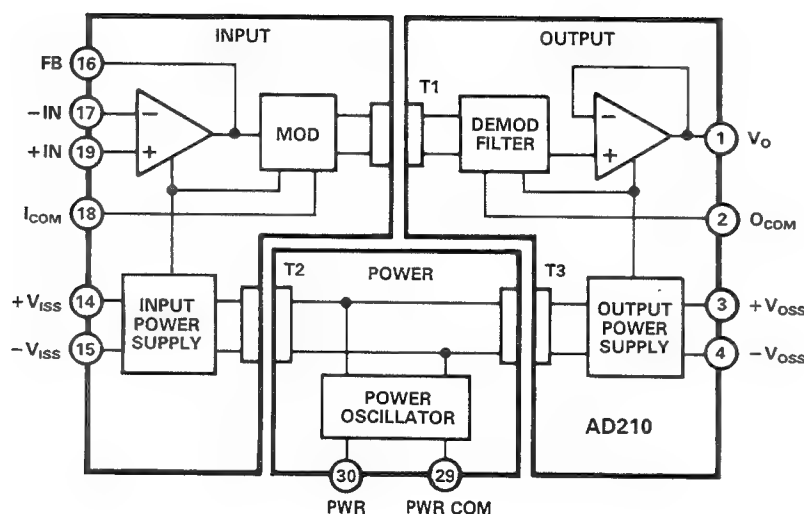


Figure 1.25

REFERENCES

1. Richard J. Higgins, **Digital Signal Processing in VLSI**, Prentice-Hall, 1990.
2. Daniel H. Sheingold, Editor, **Transducer Interfacing Handbook**, Analog Devices, Inc., 1972.
3. **High Speed Design Seminar**, Analog Devices, 1990.

SECTION II

LINEAR AND NON-LINEAR ANALOG SIGNAL PROCESSING

LINEAR AND NON-LINEAR ANALOG SIGNAL PROCESSING

- **AMPLIFIERS USED AS ANALOG SIGNAL PROCESSORS**
- **DISK DRIVE READ AMPLIFIERS**
- **ANALOG MULTIPLIERS**
- **RMS TO DC CONVERTERS**
- **LOGARITHMIC AMPLIFIERS**
- **VARIABLE GAIN AMPLIFIER (ULTRASOUND APPLICATION)**
- **PASSIVE AND ACTIVE ANALOG FILTERING**

ANTIALIASING FILTER DESIGN EXAMPLE

A PROGRAMMABLE STATE VARIABLE FILTER

SEVEN-POLE FDNR 20KHZ ANTIALIASING FILTER

WIDEBAND SALLEN-KEY FILTER

SECTION II

LINEAR AND NON-LINEAR ANALOG SIGNAL PROCESSING

AMPLIFIERS USED AS ANALOG SIGNAL PROCESSORS

2

High performance operational amplifiers can be used to perform such mathematical operations such as addition, subtraction, differentiation, and integration. They can also be used to perform non-linear operations such as rectification and exponentiation. In addition, these operations can be performed on signals having bandwidths of greater than 100MHz because of recent advances in linear IC process technology as well as circuit topologies.

An inverting adder can be made with an op amp and a few resistors as shown in Figure 2.1, or, an instrumentation amplifier can be used with another op amp to form a non-inverting adder. A simple subtractor can be made by using an op amp in the differential mode as shown in Figure 2.2, or an instrumentation amplifier can be used for high input impedance.

OP AMP AND INSTRUMENTATION AMP ADDERS

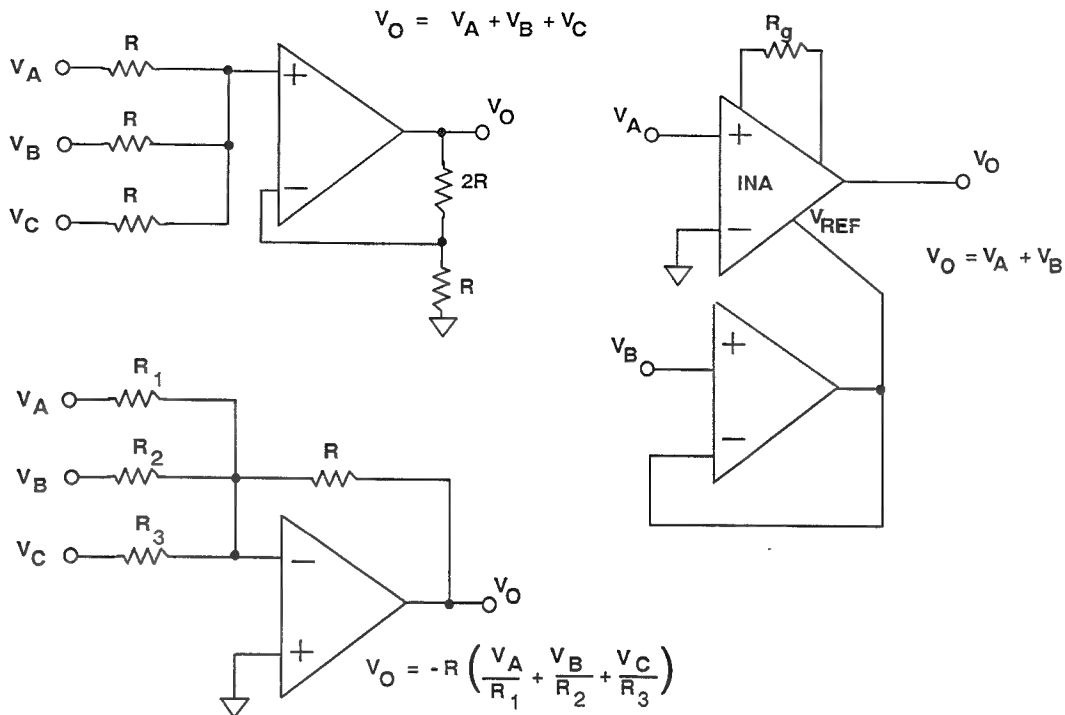


Figure 2.1

OP AMP AND INSTRUMENTATION AMP SUBTRACTORS




Figure 2.2

OP AMP DIFFERENTIATORS




Figure 2.3

OP AMP INTEGRATOR




Figure 2.4

Figure 2.3 shows two versions of a simple differentiator, and Figure 2.4 shows the classic op amp circuit used for integration.

The circuit of Figure 2.5 is a typical absolute-value circuit. It comprises a diode network and a differencing circuit. Later, we will see how a multiplier can be used to perform the same function without the problems associated with the forward voltage drops of the diodes.

The peak detector circuit of Figure 2.6 can accurately capture the amplitude of input pulses as narrow as 200ns and can hold their

value with a droop rate of less than 20μV/ms. The high bandwidth and 200V/μs slewrate of the AD843 op amp allows the detector's output to "keep up" with its input thus minimizing overshoot. The low (less than 1nA) input current of the AD843 (because of its FET input) ensures that the droop rate is limited only by the reverse leakage of diode D2, which is typically less than 10nA for the type shown. A more detailed description of the operation of this peak detector is given in Reference 1.

OP AMP ABSOLUTE VALUE CIRCUIT (FULL WAVE RECTIFIER)




Figure 2.5

A FAST PEAK DETECTOR CIRCUIT

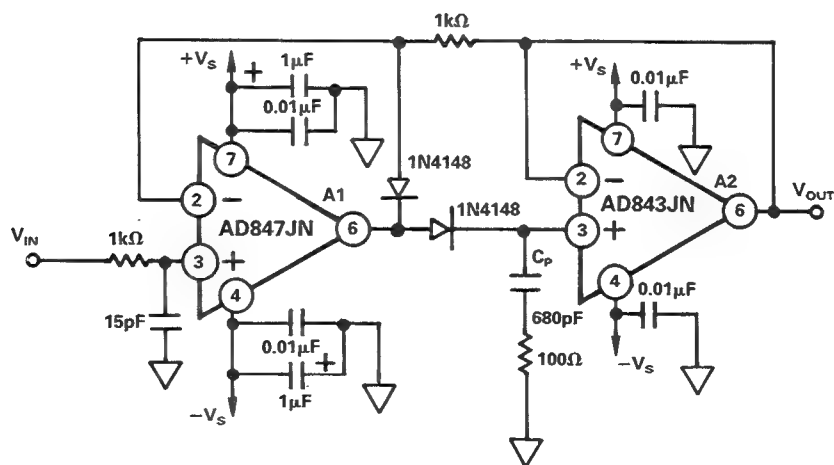


Figure 2.6

DISK DRIVE READ AMPLIFIERS

An excellent example of an analog signal processor can be found in disk drive read electronics. A block diagram of a typical system is shown in Figure 2.7. Data is stored as a stream of non-return-to-zero (NRZ) pulses, where each pulse requires a magnetic flux change on the disk. In the read mode, these flux reversals are sensed by the *head*, and a voltage is output to the preamplifier. The output from the head is a

signal varying in amplitude, and adversely affected by the noise in the system. Amplitude variations are caused by deviations in height between the head and disk media, media integrity, and electrical noise. Therefore the output of the preamplifier also varies in amplitude. Although the preamplifier output represents digital data from the disk, analog conditioning is required in order to avoid intolerable errors.

DISK DRIVE READ ELECTRONICS

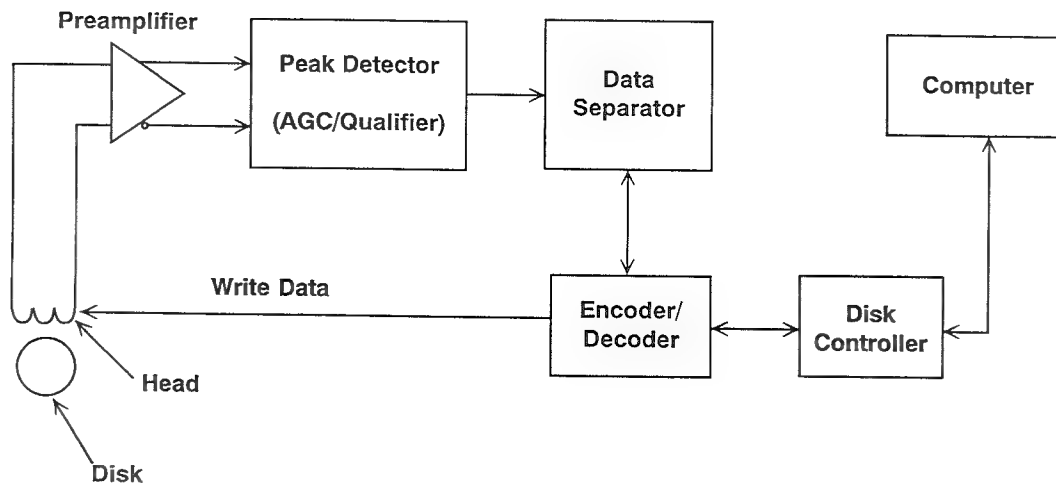


Figure 2.7

One half of the peak detection circuitry's function is to establish an automatic gain control (AGC) loop which produces a signal of constant amplitude. This may be implemented as shown in Figure 2.8 using a full-wave rectifier (FWR), a sample-and-hold (SHA) and a variable-gain-amplifier (VGA). The other half of the peak detector's function is to qualify the signal and output a digital pulse for every bit stored on the disk. Since the read signal is rectified prior to being applied to the data qualifier, a single comparator can be used to qualify both the

positive and (rectified) negative peaks. This is accomplished by verifying that the input signal has achieved a minimum amplitude, and strobing the output only when a signal peak has been detected. The AGC'd signal is differentiated to produce a zero-crossing at the signal peaks. This peak and amplitude detection scheme is necessary to minimize data errors caused by noise-induced zero-crossings. That is, the level comparators must make sure the zero-crossings did indeed occur at the signal peaks.

PEAK DETECTOR

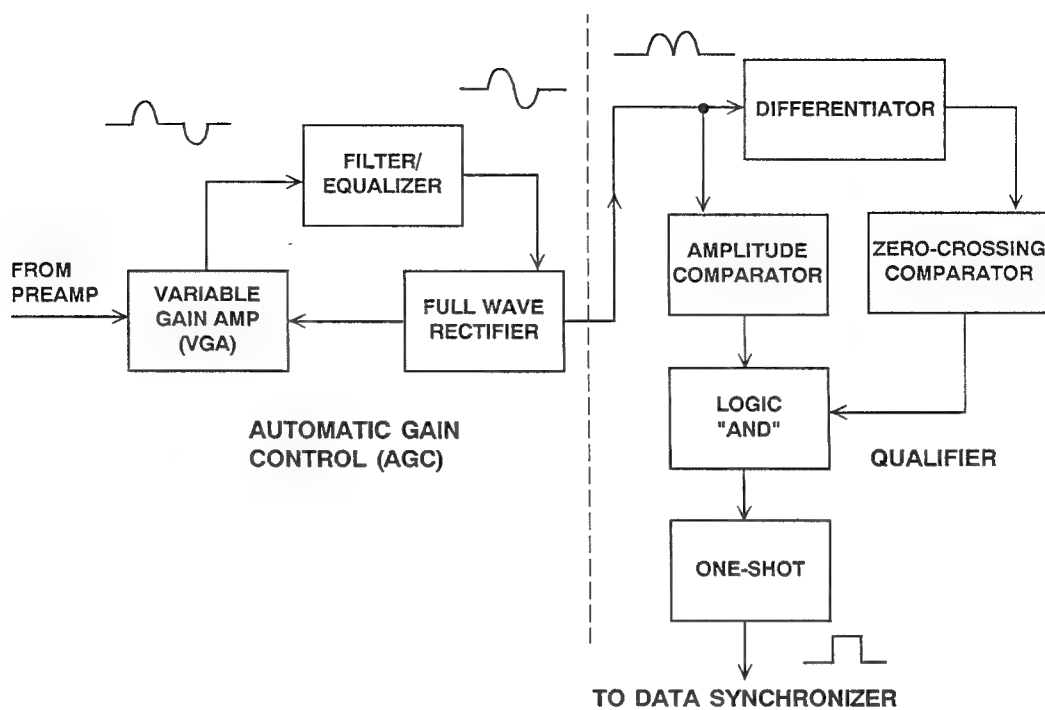


Figure 2.8

DATA SEPARATOR

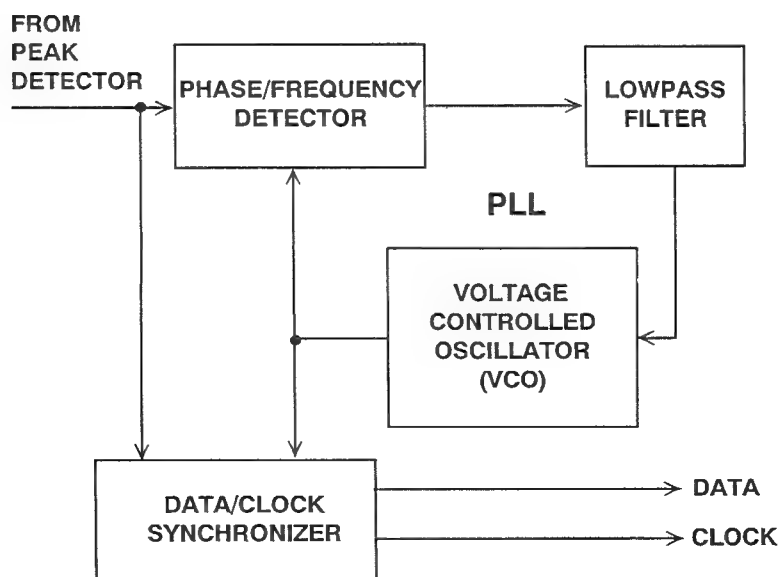


Figure 2.9

The data separator is used to extract clock information that is imbedded in the data. The input to this device is the digital pulse stream from the pulse detector's qualifier. The output is synchronized data and clock information. This is implemented using a phase-locked loop as shown in Figure 2.9.

The encoder/decoder is used for both write (encoder) and read (decoder) operations. In the write mode, NRZ data from the disk controller is coded and output to the head. In the read mode, the synchronized clock and (encoded) data outputs from the data separator are input to the encoder/decoder, whose output, in original NRZ format is passed on to the disk controller.

The AD890 (Precision Wideband Channel Processing Element) and the AD891 (Hard

Disk Data Channel Qualifier) operate together to perform the pulse detection function at a rate of up to 50Mb/s. In addition, the AD891 can be used as a stand-alone variable gain amplifier having an 80MHz bandwidth with a 30dB maximum gain and a 40dB control range.

The AD892E (ECL) and the AD892T (TTL) integrate the AGC and peak detection functions of the AD890 and the AD891 onto a single chip. The AD892E can be used up to 30Mb/s, and the AD892T up to 25Mb/s.

The AD897 is a fully integrated 40Mb/s read channel which contains the AGC, peak detector, and phase locked loop. A functional block diagram of the device is shown in Figure 2.10.

AD897 40 Mb/s DISK DRIVE READ CHANNEL

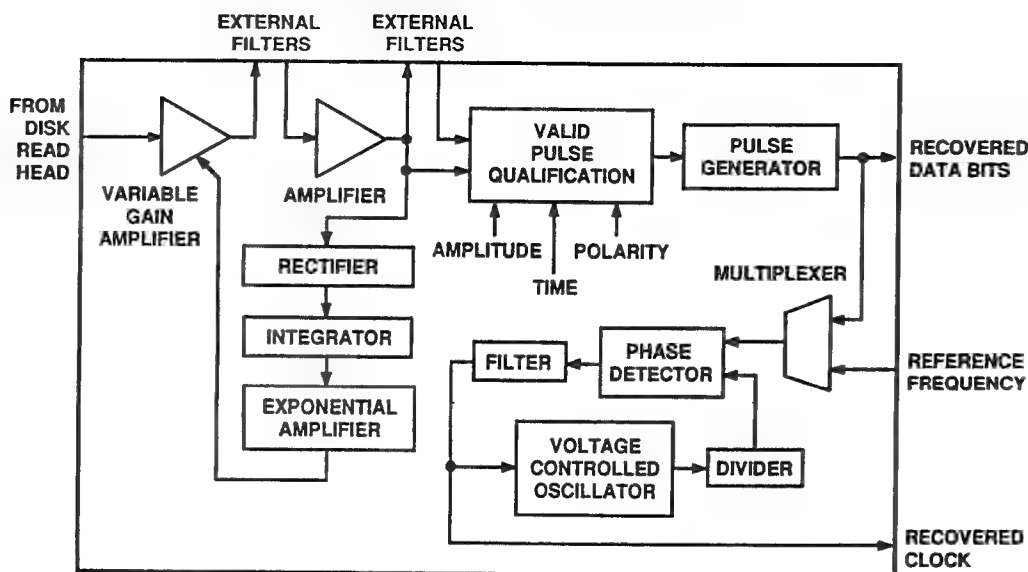


Figure 2.10

ANALOG MULTIPLIERS

A multiplier is a device having two input ports and an output port. The signal at the output is the product of the two input signals. If both input and output signals are voltages, the transfer characteristic is the product of the two voltages divided by a scaling factor, K , which has the dimension of voltage (see Figure 2.11). From a mathematical point of view, multiplication is a “four quadrant” operation - that is to say that both inputs may be either positive or negative, as may be the output. Some of the circuits used to produce electronic multipliers, however, are limited to signals of one polarity. If both

signals must be unipolar, we have a “single quadrant” multiplier, and the output will also be unipolar. If one of the signals is unipolar, but the other may have either polarity, the multiplier is a “two quadrant” multiplier, and the output may have either polarity (and is “bipolar”). The circuitry used to produce one- and two-quadrant multipliers may be simpler than that required for four quadrant multipliers, and since there are many applications where full four quadrant multiplication is not required, it is common to find accurate devices which work only in one or two quadrants.

BASIC MULTIPLIER




Figure 2.11

TYPES OF MULTIPLIERS

TYPE	V_x	V_y	V_{out}
Single Quadrant	Unipolar	Unipolar	Unipolar
Two Quadrant	Bipolar	Unipolar	Bipolar
Four Quadrant	Bipolar	Bipolar	Bipolar

Figure 2.12

The simplest electronic multipliers use logarithmic amplifiers. The computation relies on the fact that the antilog of the sum

of the logs of two numbers is the product of those numbers (see Figure 2.13).

COMPUTATION WITH LOG AND ANTILOG CIRCUITS

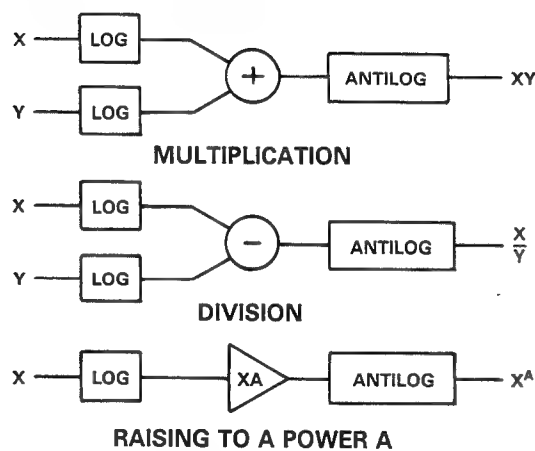


Figure 2.13

The disadvantages of this type of multiplication are the very limited bandwidth and single quadrant operation. A far better type of multiplier uses the "Gilbert Cell". This structure was invented by Barrie Gilbert in the late 1960s. (See References 2 and 3).

There is a linear relationship between the collector current of a silicon junction transistor and its transconductance (gain) which is given by

$$dI_c / dV_{be} = qI_c / kT, \text{ where}$$

I_c = the collector current

V_{be} = the base-emitter voltage

q = the electron charge (1.60219E-19)

k = Boltzmann's constant (1.38062E-23)

T = the absolute temperature.

This relationship may be exploited to construct a multiplier with a long-tailed pair of silicon transistors, as shown in Figure 2.14.

BASIC TRANSCONDUCTANCE MULTIPLIER CIRCUIT




Figure 2.14

This is a rather poor multiplier because (1) the Y input is offset by the V_{be} - which changes non-linearly with V_y ; (2) the X input is non-linear as a result of the exponential relationship between I_c and V_{be} ; and (3) the scale factor varies with temperature.

Gilbert realized that this circuit could be linearized and made temperature stable by working with currents, rather than voltages, and by exploiting the logarithmic I_c / V_{be} properties of transistors (See Figure 2.15.)

THE GILBERT CELL - A LINEAR TWO-QUADRANT MULTIPLIER

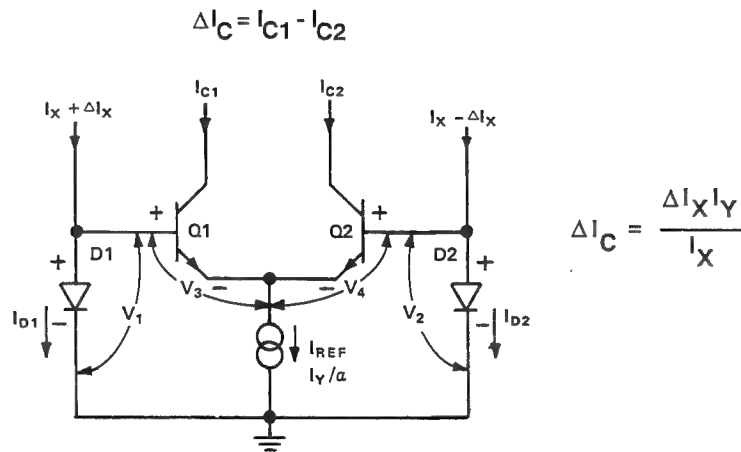


Figure 2.15

The X input to the Gilbert Cell takes the form of a differential current, and the Y input is a unipolar current. The differential X currents flow in two diode connected transistors, and the logarithmic voltages compensate for the exponential V_{be} / I_c relationship. Furthermore the q / kT scale factors cancel. This gives the Gilbert Cell the linear transfer function

$$\Delta I_c = \frac{\Delta I_x I_y}{I_x}$$

As it stands the Gilbert Cell has three inconvenient features: (1) Its X input is a differential current; (2) Its output is a differential current; and (3) Its Y input is a unipolar current - so the cell is only a two quadrant multiplier.

By cross-coupling two such cells and using two voltage-to-current converters (as shown in Figure 2.16), we can convert the basic architecture to a four quadrant device with voltage inputs, such as the AD534. At low and medium frequencies a subtractor amplifier may be used to convert the differential current at the output to a voltage. Because of its voltage output architecture, the bandwidth of the AD534 is only about 1MHz.

In Figure 2.16, Q1A & Q1B, and Q2A & Q2B form the two core long-tailed pairs of

the two Gilbert Cells, while Q3A and Q3B are the linearizing transistors for both cells. In Figure 2.16 there is an operational amplifier acting as a differential current to single-ended voltage converter, but for higher speed applications the cross-coupled collectors of Q1 and Q2 form a differential open collector current output (as in the AD834 500MHz multiplier).

The translinear multiplier relies on the matching of a number of transistors and currents. This is easily accomplished on a monolithic chip. Even the best IC processes have some residual errors, however, and these show up as four dc error terms in such multipliers (see Figure 2.17).

In early Gilbert Cell multipliers these errors had to be trimmed by means of resistors and potentiometers external to the chip, which was somewhat inconvenient. With modern analog processes, which permit the laser trimming of SiCr thin film resistors on the chip itself, it is possible to trim these errors during manufacture so that the final device has very high accuracy. Internal trimming has the additional advantage that it does not reduce the high frequency performance, as may be the case with external trim pots.

4-QUADRANT TRANSLINEAR MULTIPLIER, THE AD534

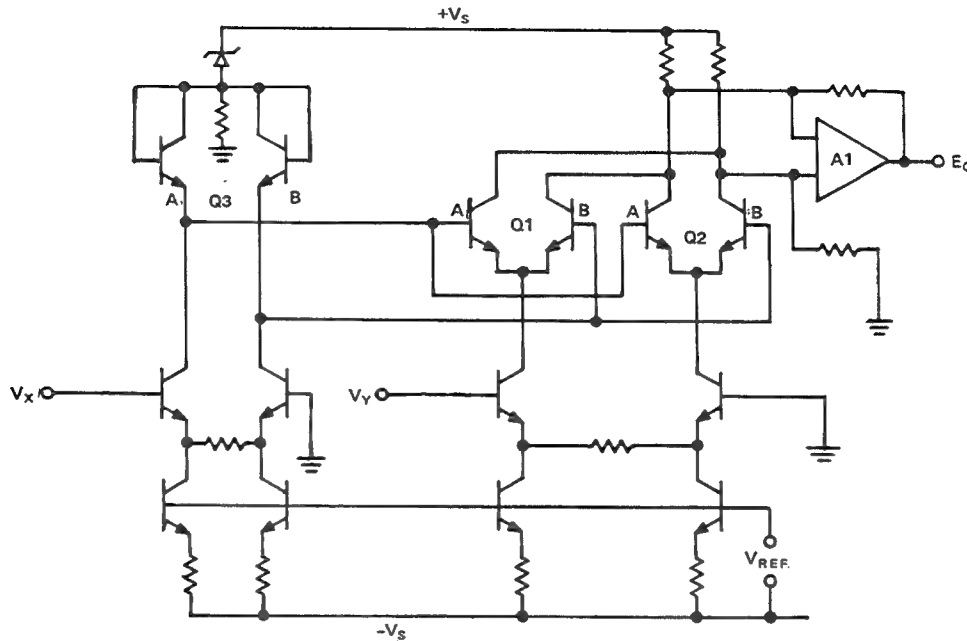


Figure 2.16

TRIMMABLE ERRORS IN MULTIPLIERS

- X-Input Offset Voltage: Y Feedthrough
- Y-Input Offset Voltage: X Feedthrough
- Z-Input (Output Amplifier)
Voltage Offset: dc Output Offset Voltage
- Resistor Mismatch: Gain Error

Figure 2.17

The AD633 is a low cost 1MHz four quadrant multiplier. The AD734 is a 10MHz precision four quadrant multiplier/divider with a voltage output, and the AD834 is a 4 quadrant multiplier with differential X

inputs, differential Y inputs, differential open collector current outputs, and a bandwidth of over 500MHz. Offset voltages are laser wafer trimmed, and accuracy is better than 0.1%.

KEY FEATURES OF THE TRANSLINEAR MULTIPLIER

- **High Accuracy:** Better than 0.1% Possible
- **Wide Bandwidth** (Over 60MHz Voltage Output, Over 500MHz Current Output)
- **Simplicity, Low Cost, and Ease of Use**

2

Figure 2.18

Multipliers can be used quite efficiently as power meters, by applying the voltage to one input and the current to the other input. Figure 2.19 shows an audio frequency power meter that measures the power output for an audio amplifier into an 8Ω load resistor. The 10kΩ-18kΩ voltage divider scales the amplifier's output swing to a maximum of 10V (from a maximum of 28V, representing about 100W peak power), well within the AD633's input range. The voltage is measured across

the load, with the tap of the divider connected to X1 and the lower end of the load to X2. Current is measured across the 0.1Ω shunt. The AD711 op amp amplifies the signal to manageable levels and to drive the multiplier's Y1 input. The output of the multiplier is $(X1 - X2)Y1/10$, which is proportional to the product of the load voltage and the load current, hence the power dissipated in the load.

AUDIO POWER METER USING AD633 MULTIPLIER

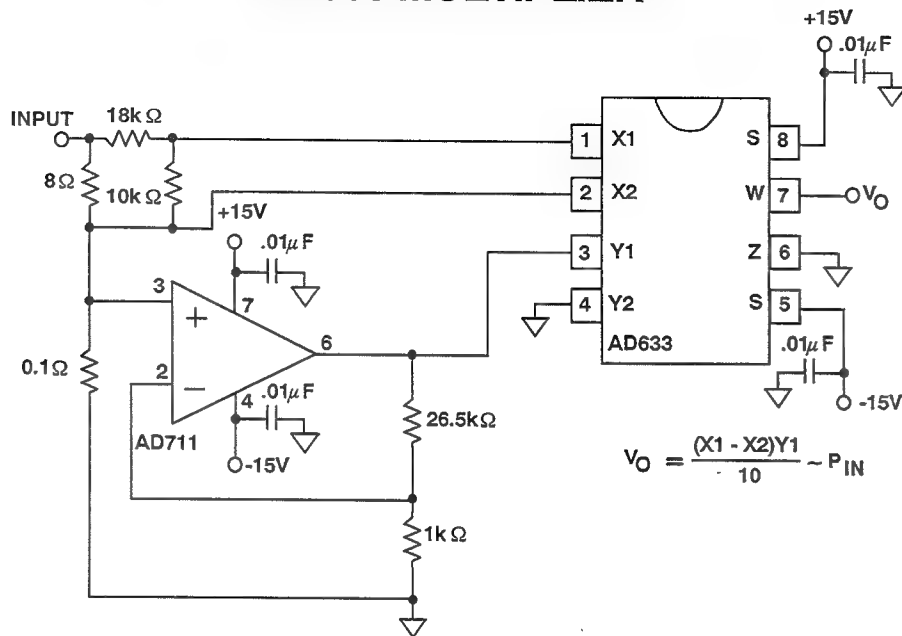


Figure 2.19

Multipliers can be placed in the feedback loop of op amps to form several useful functions. Figure 2.20 shows a multiplier and an op amp configured as a divider in both inverting and non-inverting mode. This circuit illustrates the principle of analog computa-

tion that a function generator in a negative feedback loop computes the inverse function (provided, of course, that the function is monotonic over the range of operation (see Figure 2.21).

MULTIPLIERS CONFIGURED AS DIVIDERS

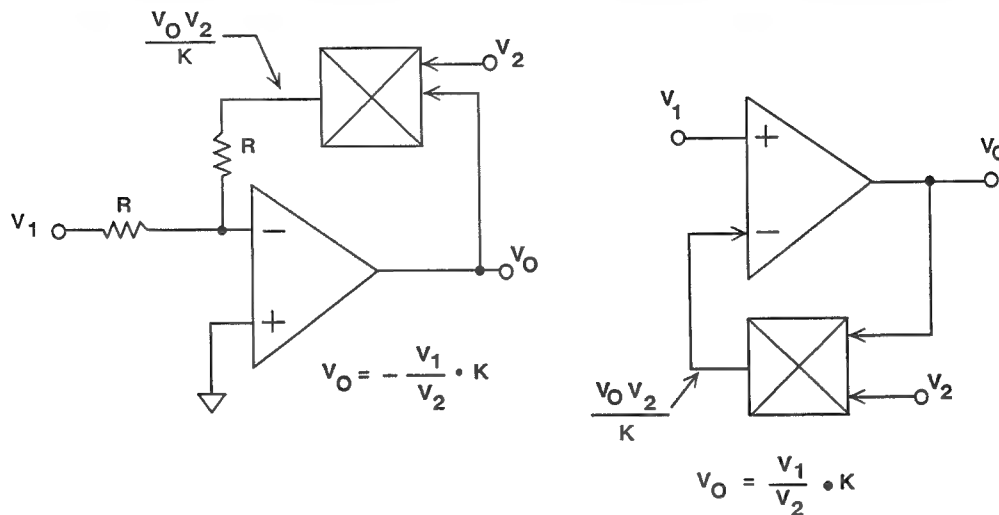



Figure 2.20

A FUNCTION GENERATOR IN A NEGATIVE-FEEDBACK LOOP GENERATES THE INVERSE FUNCTION



NOTE: FUNCTION MUST BE MONOTONIC OVER THE RELEVANT RANGE

Figure 2.21

The AD734 is a 10MHz four-quadrant multiplier with an external input to dynamically change the scale factor, thereby accomplishing a direct-divide function. A functional block diagram of the device is shown in Figure 2.22.

A modulator is closely related to a multiplier. The output of a multiplier is the instantaneous product of its inputs. The output of a modulator is the instantaneous product of a signal on one of its inputs (known as the signal input) and the *sign* of the signal on the other input (known as the carrier input). A modulator may be modelled as an amplifier whose gain is switched positive and negative by the output of a comparator on its carrier input (as in the case of the AD630 balanced modulator) - or as a multiplier with a high-gain limiting amplifier between the carrier output and one of its ports. Both configurations are shown in Figure 2.23. Most high speed integrated

circuit modulators consist of the translinear multiplier with a limiting amplifier in the carrier path. A precision rectifier, or absolute value circuit, can be made using a modulator as shown in Figure 2.24.

A multiplier can be used as a variable-gain amplifier as shown in Figure 2.25. The control voltage is applied to one input, and the signal to the other. The AD539 two-quadrant multiplier (60MHz bandwidth) and the AD844 current feedback amplifier can be used in a 20MHz variable gain amplifier configuration as shown in Figure 2.26. The frequency response of the variable gain amplifier for gains of +4 to -46dB as well as the transient response is shown in Figure 2.27. A current feedback (or transimpedance) amplifier is ideally suited for this application since its bandwidth remains relatively constant over a wide range of closed-loop gains.

10 MHz MULTIPLIER WITH DIRECT DIVIDE CAPABILITY

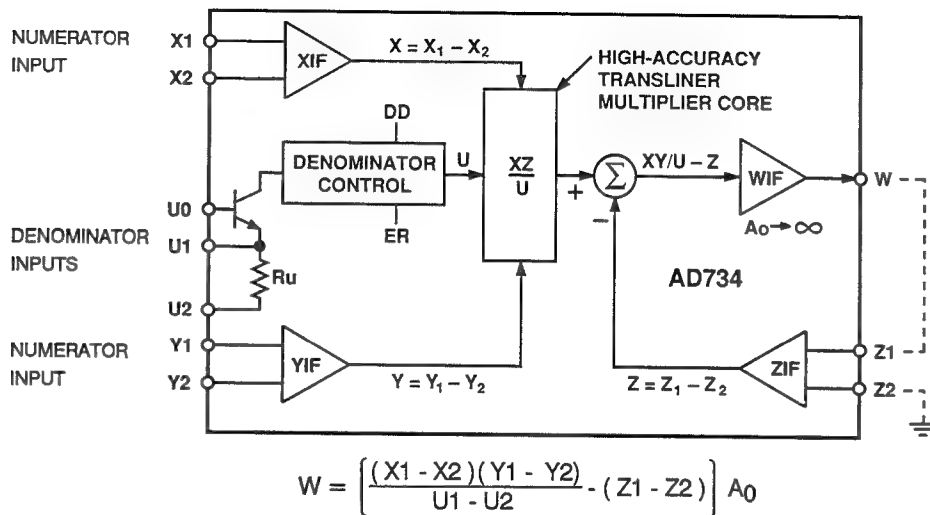


Figure 2.22

TWO MODULATOR MODELS

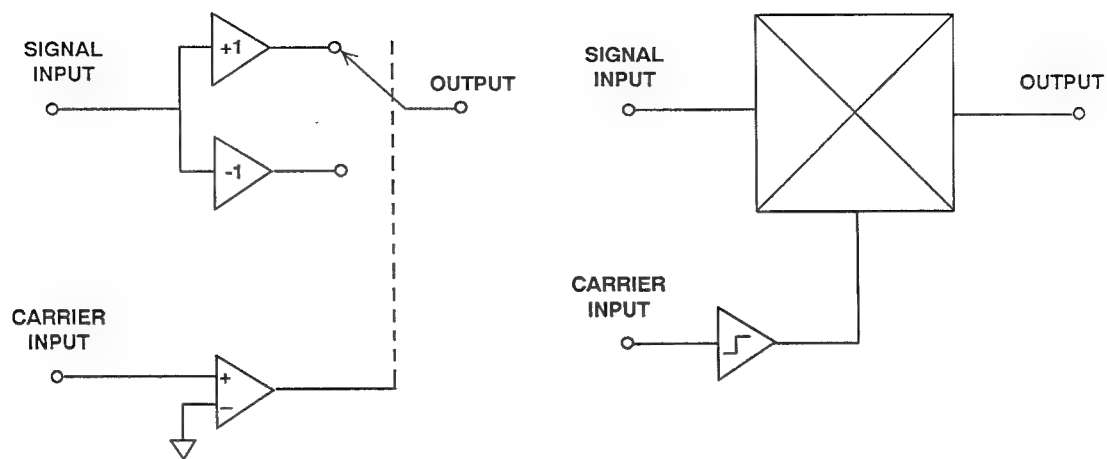


Figure 2.23

MODULATOR AS A PRECISION RECTIFIER (ABSOLUTE VALUE CIRCUIT)




Figure 2.24

MULTIPLIER AS A VARIABLE GAIN AMPLIFIER

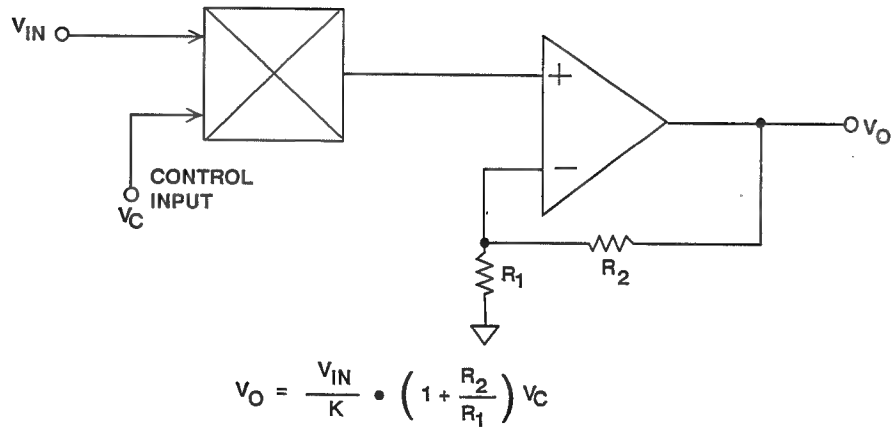


Figure 2.25

20MHz VARIABLE GAIN AMPLIFIER USING THE AD539

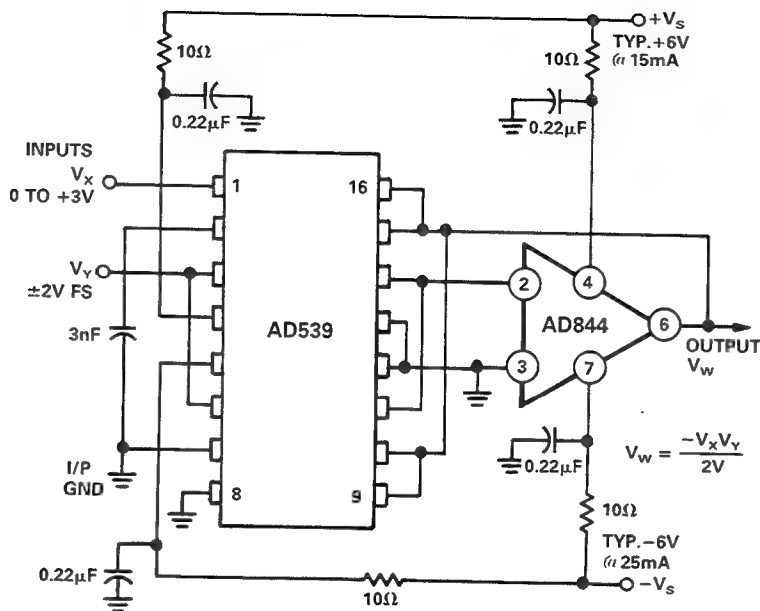




Figure 2.26

VGA AC AND TRANSIENT RESPONSE



VGA ac Response



VGA Transient Response with $V_x = 1V, 2V, \text{ and } 3V$

Figure 2.27

RMS TO DC CONVERTERS

The rms or root mean square is a fundamental measurement of the magnitude of an ac signal. Defined practically, the rms value assigned to the ac signal is the amount of dc required to produce an equivalent amount of heat in the same load. Defined mathematically, the rms value of a voltage is defined as the value obtained by squaring the signal, taking the average, and then taking the square root. The averaging time must be sufficiently long to allow filtering at the lowest frequencies of operation desired. A complete discussion of rms to dc converters can be found in Reference 5, but we will now show a few examples of how efficiently analog circuits discussed so far can perform this function.

The first method, called the explicit method, is shown in Figure 2.28. The input signal is first squared by a multiplier. The average value is then taken by using an appropriate filter, and the square root is taken using an op amp with a second multi-

plier in the feedback loop. This circuit has limited dynamic range because the stages following the squarer must try to deal with a signal that varies enormously in amplitude. This restricts this method to inputs which have a maximum dynamic range of approximately 10:1 (20dB). However, excellent bandwidth (greater than 100MHz) can be achieved with high accuracy if a multiplier such as the AD834 is used as a building block (see Figure 2.29).

Figure 2.30 shows the circuit for computing the rms value of a signal using the implicit method. Here, the output is fed back to the direct-divide input of a multiplier such as the AD734. In this circuit, the output of the multiplier varies linearly (instead of as the square) with the rms value of the input. This considerably increases the dynamic range of the implicit circuit as compared to the explicit circuit. The disadvantage of this approach is that it generally has less bandwidth than the explicit computation.

EXPLICIT RMS COMPUTATION

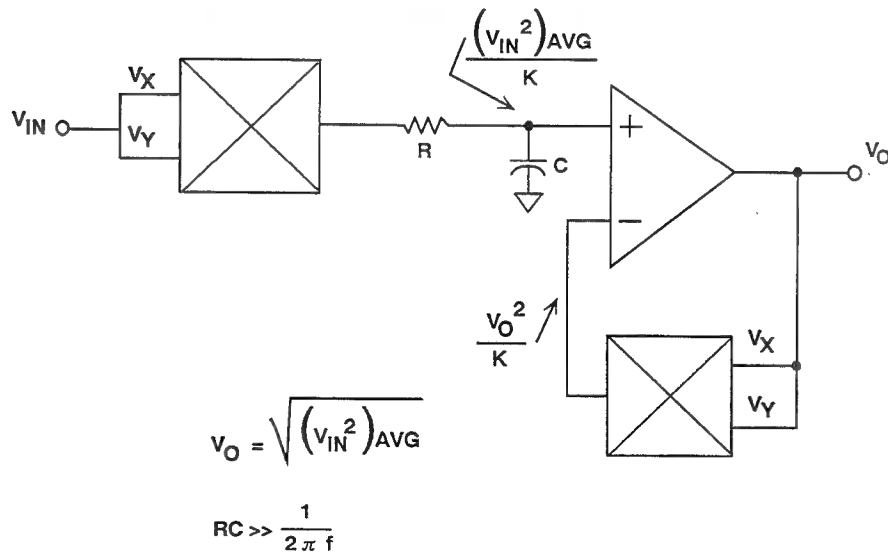


Figure 2.28

WIDEBAND R.M.S. MEASUREMENT WITH THE AD834

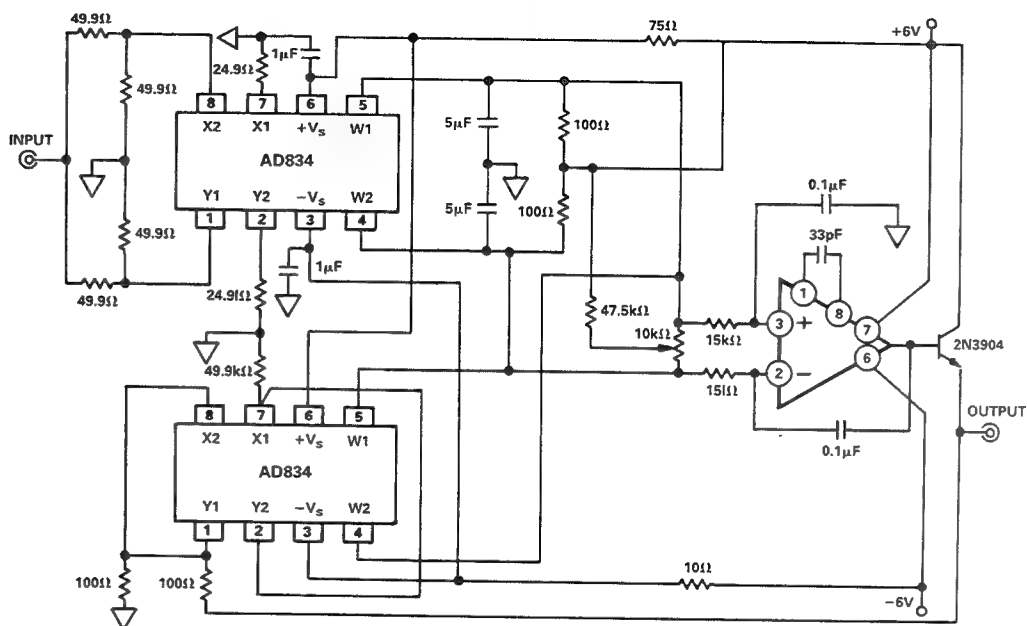


Figure 2.29

IMPLICIT RMS COMPUTATION

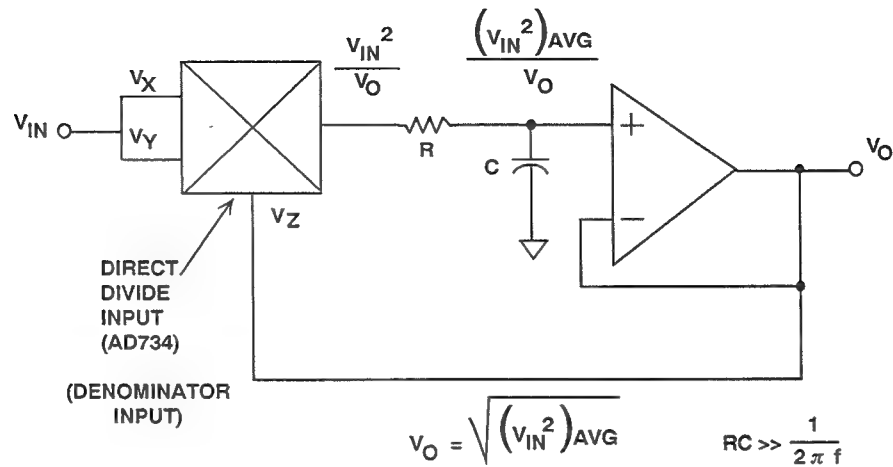


Figure 2.30

LOGARITHMIC CONVERTERS

Log amps find wide applications where signals having wide dynamic ranges (greater than 100dB, perhaps) must be processed by elements such as ADCs which may have more limited dynamic ranges. Log amps have maximum incremental gain for small signals; the gain decreases in inverse proportion to the magnitude of the input. This permits the amplifier to accept signals with a wide input dynamic range and compress them substantially.

The term "Logarithmic Amplifier" (generally abbreviated to "log amp") is something of a misnomer, and "logarithmic converter" would be a better description. A log amp must satisfy a transfer function of the form

$$V_{out} = V_y \log(V_{in}/V_x)$$

over some range of input values which may vary from 100:1 (40dB) to over 1,000,000:1 (120dB).

With inputs very close to zero, log amps cease to behave logarithmically, and most

then have a linear V_{in}/V_{out} law. This behavior is often lost in device noise. Noise often limits the dynamic range of a log amp. The constant V_y has the dimensions of voltage because the output is a voltage. The input, V_{in} , is divided by a voltage V_x , because the argument of a logarithm must be a simple dimensionless ratio.

A graph of the transfer characteristic of a log amp is shown in Figure 2.31. The scale of the horizontal axis (the input) is logarithmic, and the ideal transfer characteristic is a straight line. When $V_{in} = V_x$, the logarithm is zero ($\log 1 = 0$). V_x is therefore known as the *intercept voltage* of the log amp because the graph crosses the horizontal axis at this value of V_{in} .

The slope of the line is proportional to V_y . When setting scales, logarithms to the base 10 are most often used because this simplifies the relationship to decibel values: when $V_{in} = 10 V_x$ the logarithm has the value of 1, so the output voltage is V_y . When $V_{in} = 100 V_x$ the output is $2V_y$ and so forth. V_y can

LOG AMP TRANSFER FUNCTION




Figure 2.31

therefore be viewed either as the “slope voltage” or as the “volts per decade factor.”

Log amps can respond to negative inputs in three different ways: (1) They can give a fullscale negative output as shown in Figure 2.32. (2) They can give an output which is proportional to the log of the absolute value of the input and disregards its sign as shown in Figure 2.33. This type of log amp can be considered to be a full-wave detector with a logarithmic characteristic, and is often referred to as a detecting log amp. (3) They can give an output which is proportional to the log of the absolute value of the input and has the same sign as the input as shown in Figure 2.34. This type of log amp can be considered to be a video amp with a logarithmic characteristic, and may be known as a logarithmic video amplifier or, sometimes, a “true log amp.”

There are three basic architectures which may be used to produce log amps: the basic diode log amp, the successive detection log amp, and the “true log amp” which is based on cascaded semi-limiting amplifiers.

The voltage across a silicon diode is pro-

portional to the logarithm of the current through it. If a diode is placed in the feedback path of an inverting op-amp, the output voltage will be proportional to the log of the input current as shown in Figure 2.35. In practice, the dynamic range of this configuration is limited to 40-60dB because of non-ideal diode characteristic, but if the diode is replaced with a diode-connected transistor as shown in Figure 2.36, the dynamic range can be extended to 120dB or more. This circuit has several major drawbacks; external compensation networks are required in order to obtain reasonable dc performance, and frequency response is limited to a few hundred kHz because of the feedback capacitance. Several such log amps can be combined on a single chip to perform an analog computer which performs both log and anti-log operations as in the AD538 shown in Figure 2.37. The temperature variation in the log operations is unimportant, since it is compensated by similar variation in the anti-logging. The AD538 can multiply, divide, and raise to powers.

BASIC LOG AMP (SATURATES NEGATIVE WITH NEGATIVE INPUT)




Figure 2.32

DETECTING LOG AMP (OUTPUT POLARITY INDEPENDENT OF INPUT POLARITY)




Figure 2.33

LOG VIDEO AMP OR "TRUE LOG AMP" (SYMMETRICAL RESPONSE TO POSITIVE OR NEGATIVE SIGNALS)




Figure 2.34

THE DIODE/OP-AMP LOG AMP




Figure 2.35

TRANSISTOR LOG AMP


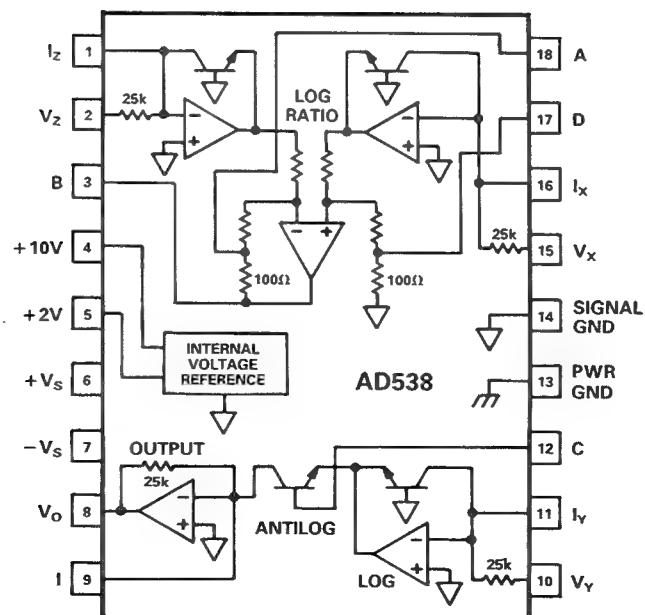


Figure 2.36

THE AD538



$$V_{out} = V_y \left(\frac{V_z}{V_x} \right)^m$$

Figure 2.37

For high frequency applications, therefore, detecting and “true log” architectures are used. Although these differ in detail, the general principle behind their design is common to both: instead of one amplifier having a logarithmic characteristic, these designs use a number of similar cascaded linear stages having well-defined large signal behavior.

Consider N cascaded limiting amplifiers, the output of each driving a summing circuit as well as the next stage as shown in Figure 2.38. If each amplifier has a gain of A dB, the small signal gain of the strip is NA dB. If the input signal is small enough for the last stage not to limit, the output of the summing amplifier will be dominated by the output of the last stage.

2

BASIC MULTI-STAGE LOG AMP ARCHITECTURE




Figure 2.38

As the input signal increases, the last stage will limit. It will now make a fixed contribution to the output of the summing amplifier, but the incremental gain to the summing amplifier will drop to $(N-1)A$ dB. As the input continues to increase, this stage in turn will limit and make a fixed contribution to the output, and the incremental gain will drop to $(N-2)A$ dB, and so forth - until the first stage limits, and the output ceases to change with increasing signal input.

The response curve is thus a set of straight lines as shown in Figure 2.39. The total of these lines, though, is a very good approximation to a logarithmic curve, and in practical cases is an even better one, because few limiting amplifiers, especially high frequency ones, limit quite as abruptly as this model assumes.

The choice of gain, A , will also affect the log linearity. If the gain is too high the log approximation will be poor, if it is too low, too many stages will be required to achieve the desired dynamic range. Generally, gains of 10 to 12dB ($3\times$ to $4\times$) are chosen.

The specifications for log amps will, of course, include noise, dynamic range, frequency response (some of the amplifiers used as successive detection log amp stages have low frequency as well as high frequency cutoff), the slope of the transfer characteristic (expressed in V/dB or mA/dB depending on whether we are considering a voltage - or current output device), the intercept point (the input level at which the output voltage or current is zero), and the log linearity. (See Figure 2.40).

BASIC MULTI-STAGE LOG AMP RESPONSE (UNIPOLAR CASE)




Figure 2.39

KEY PARAMETERS OF LOG AMPS

- **NOISE**
The Noise Referred to the Input (RTI) of the Log Amp.
It May Be Expressed as a Noise Figure or as a Noise Spectral Density (Voltage, Current, or Both) or as a Noise Voltage, a Noise Current, or Both
- **DYNAMIC RANGE**
Range of Signal Over Which the Amplifier Behaves in a Logarithmic Manner (Expressed in dB)
- **FREQUENCY RESPONSE**
Range of Frequencies Over Which the Log Amp Functions Correctly
- **SLOPE**
Gradient of Transfer Characteristic in V/dB or mA/dB
- **INTERCEPT POINT**
Value of Input Signal at Which Output is Zero
- **LOG LINEARITY**
Deviation of Transfer Characteristic (Plotted on log/lin Axes) from a Straight Line (Expressed in dB)


2

Figure 2.40

The successive detection log amp consists of cascaded limiting stages as described above, but instead of summing their outputs directly, these outputs are applied to detectors, and the detector outputs are summed as shown in Figure 2.41. If the detectors have

current outputs, the summing process may involve no more than connecting all the detector outputs together. The log output of a successive detection log amp generally contains amplitude information, and the phase and frequency information is lost.

SUCCESSIVE DETECTION LOGARITHMIC AMPLIFIER



Detectors may be full- or half-wave but should be current output devices (not simple diodes) so that the detector outputs may be summed without additional summing components being necessary.

Figure 2.41

In the past, it has been necessary to construct high performance, high frequency successive detection log amps (called log strips) using a number of individual limiting amplifiers. These are typically assembled in complex and costly hybrids. Recent advances in IC processes have allowed this complete function to be integrated into a single chip.

The AD640 log amp contains five limiting stages (10dB per stage) and five full-wave detectors in a single IC package, and its logarithmic

performance extends from dc to 145MHz. Furthermore, its amplifier and full-wave detector stages are balanced so that, with reasonable well-considered layout, instability from feedback via supply rails is unlikely. A block diagram of the AD640 is shown in Figure 2.42. Unlike all previous integrated circuit log amps, the AD640 is laser trimmed to high absolute accuracy of both slope and intercept, and is fully temperature compensated. Key features of the AD640 are summarized in Figure 2.43.

BLOCK DIAGRAM OF THE AD640

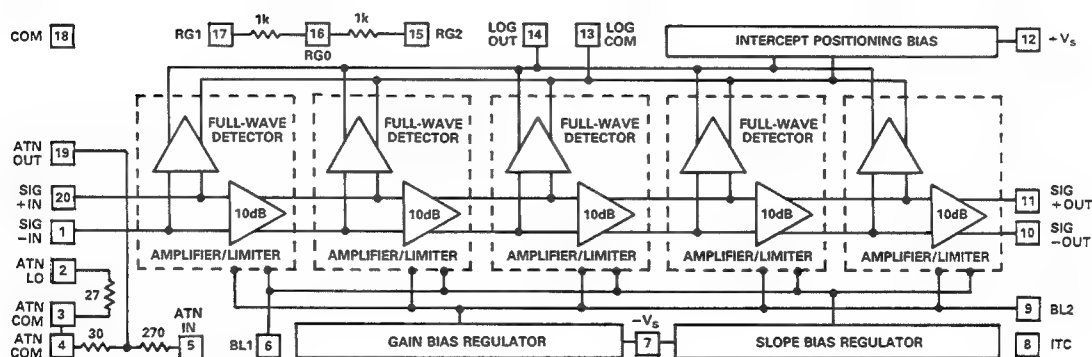


Figure 2.42

AD640 FEATURES

- 45dB Dynamic Range - Two AD640s Cascadable to 95dB
- Bandwidth dc to 145MHz - 120MHz when Cascaded
- Laser-Trimmed Slope of 1mA/decade - Temperature Stable
- Less than 1dB Log Non-Linearity
- Balanced Circuitry for Stability
- Minimal External Component Requirement

2

Figure 2.43

Each of the five stages in the AD640 has a gain of 10dB and a full-wave detected output. The transfer function for the device is shown in Figure 2.44 along with the error curve. Note the excellent log linearity over an input range of 1 to 100mV (40dB). Although well

suited to rf applications, the AD640 is dc-coupled throughout. This allows it to be used in LF and VLF systems, including audio measurements, sonar, and other instrumentation applications requiring operation to low frequencies or even dc.

DC LOGARITHMIC TRANSFER FUNCTION AND ERROR CURVE FOR SINGLE AD640




Figure 2.44

When two AD640s are cascaded, the second will be delivering an output from the noise of the first. If the full potential dynamic range is to be realized, the bandwidth must be limited because of noise. This may be done with high-pass, low-pass, or band-pass filters, depending on the required response, but, the voltage gain of these filters

in their passband must be unity, or there will be a kink in the log response. Figure 2.45 shows a 70dB log amp for broadband operation from 50 to 150MHz. The 100MHz passband limits the possible dynamic range, but the performance is still exceptional. Figure 2.46 shows a 95dB 10Hz to 100kHz log amp using two cascaded AD640s.

70 dB LOG AMP FOR 50-150 MHz USING TWO AD640

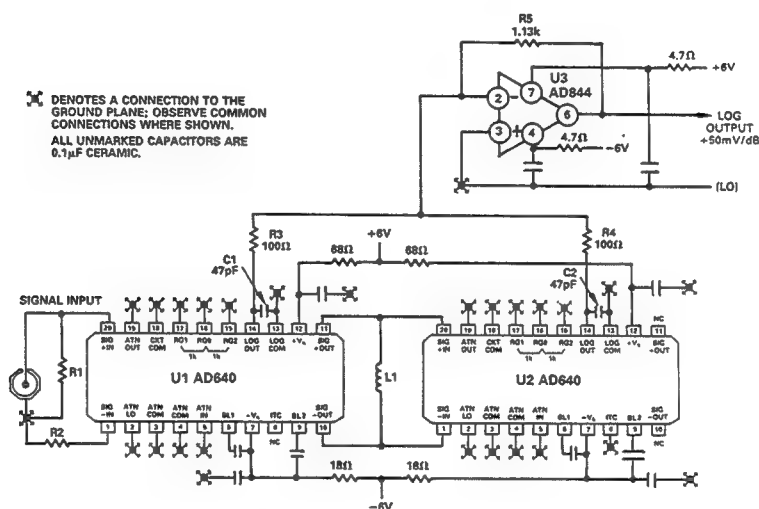


Figure 2.45

95 dB L.F. LOG AMP (10Hz - 100kHz) USING TWO AD640

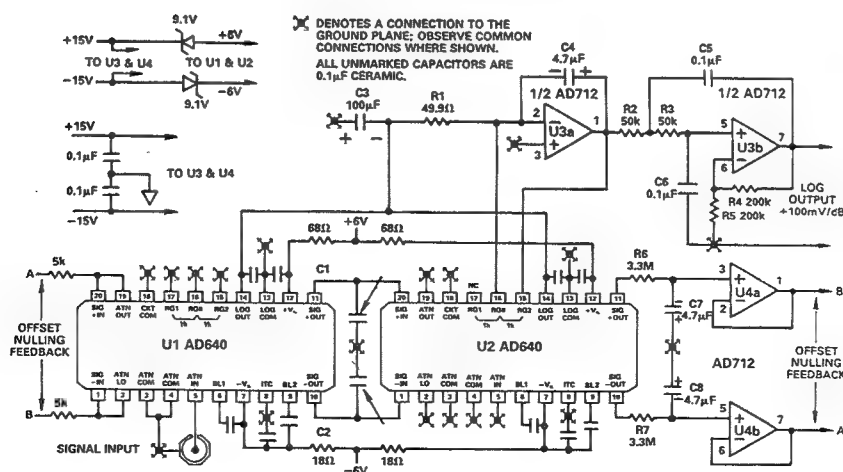


Figure 2.46

An external op amp can be used to convert the AD640 output current to a buffered output voltage as shown in Figure 2.47. The input to the AD640 (1mV to 100mV, or 40dB) results in an op amp output of 0 to 2V. If this output is applied to an 8 bit flash ADC having a corresponding input range, the weight of the ADC least significant bit (LSB) is 0.157dB reflected to the input of the AD640. For input signals near zero, the LSB value (reflected to the AD640 input) is approximately 0.02mV, while for signals approaching 100mV, the LSB value is approximately 2mV. This corresponds to an effective dynamic range of $20\log_{10}(100\text{mV}/0.02\text{mV})$, or

74dB. The 50dB dynamic range of the 8 bit flash converter has therefore been increased to 74dB (equivalent to a 12bit ADC) through the use of the 40dB AD640 and the op amp.

A “true log”, or log video amp is made up of individual cascaded gain stages, each having a small signal gain of A and a large signal (incremental) gain of unity (0dB). Each stage can be modelled as shown in Figure 2.48 as two parallel amplifiers, a limiting one with gain, and a unity gain buffer, which together feed a summing amplifier. These stages, when cascaded, form a log amp without the necessity of summing from the individual stages.

AD640 DRIVING AD770 FLASH CONVERTER

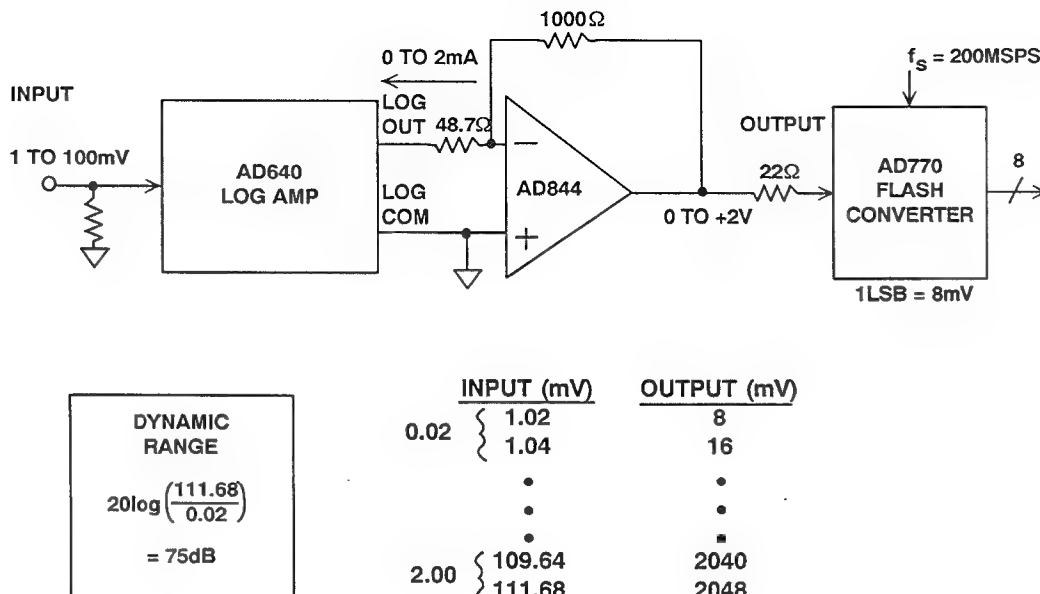


Figure 2.47

STRUCTURE AND PERFORMANCE OF TRUE LOG AMPLIFIER ELEMENT AND OF A LOG AMP FORMED BY CASCADING SEVERAL SUCH ELEMENTS

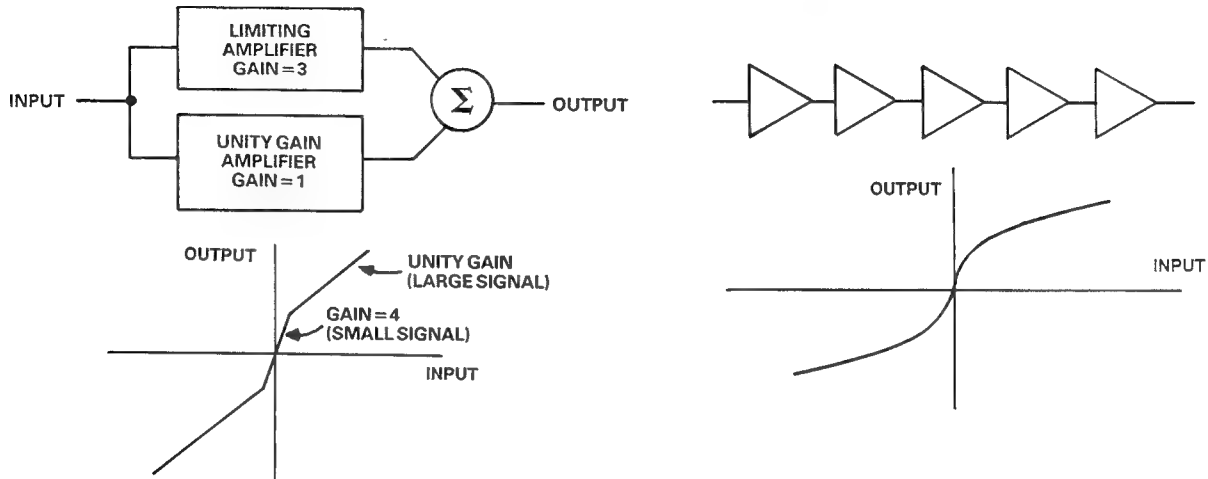


Figure 2.48

VARIABLE GAIN AMPLIFIER (ULTRASOUND APPLICATION)

A block diagram of a typical ultrasound system is shown in Figure 2.49. A burst of ultrasound energy (1 to 13MHz) is generated in an electromechanical piezoelectric transducer which physically contacts the outer body surface. The velocity of propagation of the ultrasound waves in most soft-body tissues (air and bones are the exception) is about 1500m/sec. Echoes are produced at interfaces between various types of soft-body structures. The round-trip time of each echo is used to determine its distance from the transducer.

Soft-body tissues attenuate the burst of ultrasound energy by approximately 1dB/cm/MHz. For the thicker parts of the body, as in abdominal imaging, frequencies of 1 to 2MHz are common. For imaging of shorter path lengths, as in studies of the eye or other superficial structures, frequencies as high as 20MHz can be used. Because of soft-body tissue attenuation, the receiving transducer will see a dynamic range of 100dB when

scanning from 1 to 10cm at 10MHz, independent of the tissue variations that need to be observed. Add the 50dB dynamic range typical for variations in tissue, and the transducer must have close to 150dB dynamic range. For this reason, the transducer output is usually applied to a Time Gain Amplifier (TGA) whose gain in dB is directly proportional to the amount of time elapsed from the transmission of the burst (see Figure 2.50).

For reflections that are near the surface, there will be little attenuation. For deep signal returns, gain is applied to compensate for the path attenuation. The TGA thus compensates for normal signal attenuation associated with delay/distance. The receiver (ADC) therefore only sees the intensity variations associated with the different tissue types. In scans where the propagation path is primarily soft tissue structures of comparable attenuation, such as the abdomen, a fixed gain versus time function is

B-SCAN ULTRASOUND SYSTEM BLOCK DIAGRAM




Figure 2.49

TIME GAIN AMPLIFIER

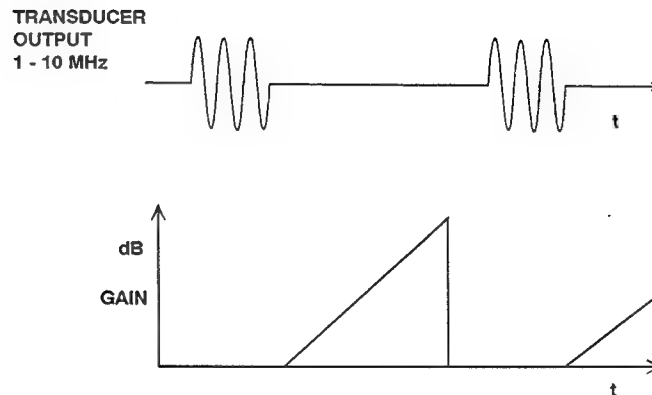


Figure 2.50

usually adequate. In other cases involving blood pools or fixed regions, it is often desirable to vary the gain versus time function. Many commercial systems make this option available. In some cases it is even desirable for the operator to calibrate the TGA on a per-patient basis in order to achieve the best diagnostic image.

In phased array ultrasound systems, the angular information is precisely determined by phasing the delays from a number of transducers (transmitted and received) to electronically select the angle to be processed. The first generation of phased array elements used analog beam forming techniques as shown in Figure 2.51. Delays at the transmitter and receiver are adjusted using variable delay filters. The next gen-

eration of phased arrays will be digital. Low cost, low power, high performance ADCs and DSPs make it practical to digitize the rf directly and digitally control the delay requirements as shown in Figure 2.52. This technique is often referred to as digital beamforming and is also being considered for use in radar systems.

The AD600 is a low-noise, variable gain amplifier optimized for use in ultrasound systems as the TGA. A block diagram is shown in Figure 2.53, and key specifications are given in Figure 2.54. The AD600 can be configured either as a dual channel device, with each channel providing 40dB of gain range, or as a single channel device providing 80dB of gain range.

ANALOG BEAMFORMING

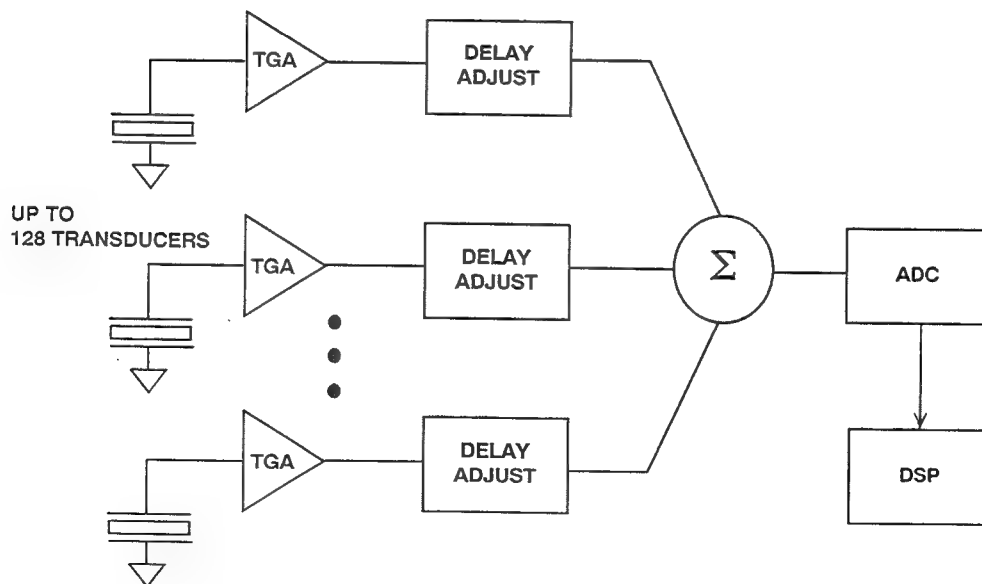


Figure 2.51

DIGITAL BEAMFORMING

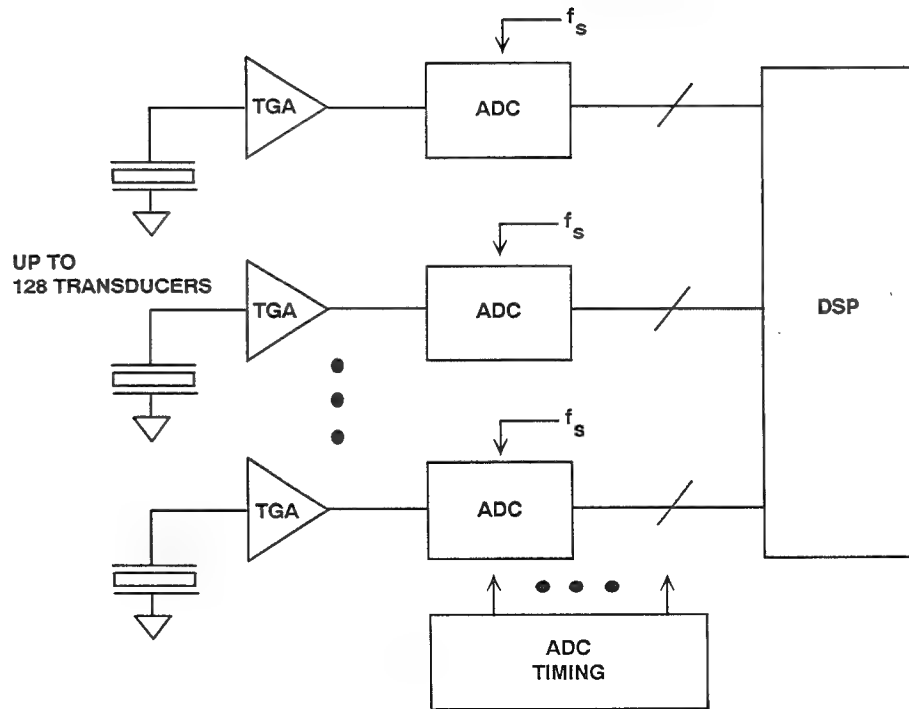


Figure 2.52

AD600 VARIABLE GAIN LOW NOISE WIDEBAND AMPLIFIER




Figure 2.53

AD600 VARIABLE GAIN AMPLIFIER KEY SPECIFICATIONS

- Two independently controlled variable gain amplifiers
- Dual channel configuration with 40dB gain control range
- Single channel configuration with 80dB gain control range
- Signal gating for each channel
- 30MHz 3dB bandwidth
- Low Distortion: -60dB THD
- 32dB per volt scaling factor
- $\pm 2.5V$ output into 200Ω load
- $\pm 5V$ supplies, 125mW power consumption
- Ideal for Ultrasound applications

Figure 2.54

Each variable gain amplifier comprises a low distortion fixed-gain (41dB) feedback amplifier, preceded by a voltage-controlled attenuator (0 to -42dB), and gain-control circuitry for smoothly interpolating the attenuator. The differential high impedance control inputs have a scale factor of 32dB/V. In addition to the gain-control inputs, each amplifier has an independent gating function which blocks transmission and sets the amplifier's dc output level to within a few millivolts of the output ground. An internal voltage reference is used to calibrate all scaling parameters. The 3dB bandwidth of each variable amplifier is nominally 35MHz and is essentially independent of gain setting. Typical input noise spectral density is $1.4nV/\sqrt{Hz}$. Signal to noise ratio in a 1MHz bandwidth is 76dB, and total harmonic

distortion is greater than 60dB.

Figure 2.55 shows the AD600 interfaced with the AD9060 (10 bits, 60MSPS) flash converter. The gain-control input is a linear ramp which is generated at the appropriate time with respect to the transmission. This example clearly demonstrates the power of the combination of analog and digital signal processing to solve a system problem. In order to provide the same dynamic range without the TGA (i.e. the ADC interfaces directly to the transducer via a fixed-gain preamp), the ADC would have to have a dynamic range of approximately 100dB. This implies a 16 bit ADC which would have to operate at a sampling frequency of at least 30MSPS — a requirement which is clearly beyond the present state of the art in ADC technology!

SINGLE CHANNEL AD600/AD9060 INTERFACE FOR ULTRASOUND

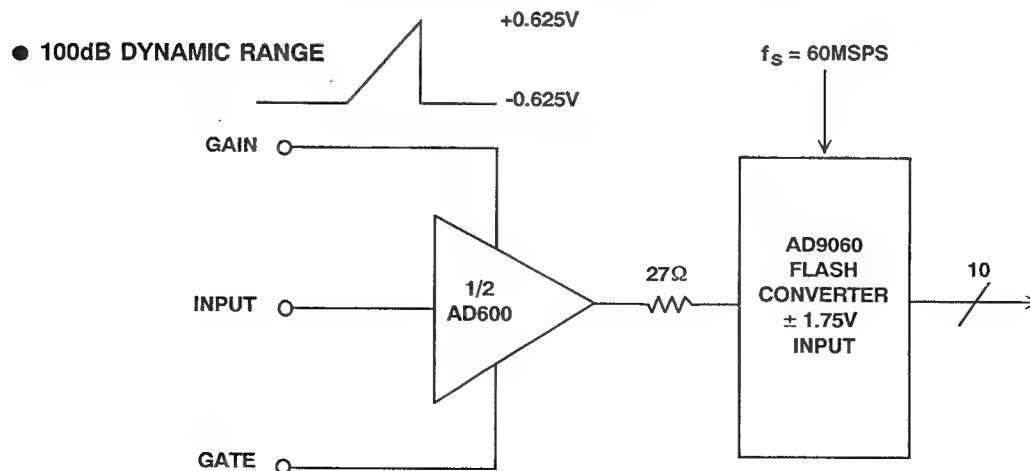


Figure 2.55

PASSIVE AND ACTIVE ANALOG FILTERING

Filtering is an important part of analog signal processing. Filtering can be used to reduce unwanted signals, limit bandwidth, help recover wanted signals, minimize aliasing in sampled data systems, and smooth the output of DACs. There are five classes of filters. *Lowpass* filters pass all frequencies below the cutoff frequency and block all frequencies above the cutoff frequency. *Highpass* filters are the inverse of the low-

pass filters. They block the low frequencies and pass those above the cutoff frequency. *Bandpass* filters pass those frequencies between the lower cutoff and upper cutoff frequencies and reject all others. *Bandstop* filters are the inverse of bandpass filters. They reject frequencies between the cutoff frequencies and pass all others. *Allpass* filters pass all frequencies equally but introduce a predictable phase delay to the signal.

CLASSES OF PASSIVE AND FILTERS

- Lowpass
- Highpass
- Bandpass
- Bandstop
- Allpass

Figure 2.56

Traditional filters were passive, that is designed with no active elements. Active components were too costly and had very poor performance characteristics. Inductors, capacitors, and resistors were used to synthesize the filter. This approach has several difficulties because inductors become physically large for low frequency filters and have poor characteristics at high frequencies.

There is a great deal of interaction between the different sections of the filter. Impedance levels must be precisely controlled. Close component tolerances are difficult to manufacture and maintain. Despite these limitations passive filters are still dominant at high frequencies, primarily due to dynamic performance limitations of op amps.

PASSIVE FILTERS

- Designed with Inductors, Capacitors, Resistors
- Large Inductors Required for Low Frequency Filters
- Interaction Between Filter Stages
- Component Tolerances Difficult to Manufacture and Maintain
- Still the Only Solution at High Frequencies Due to Active Component Limitations

Figure 2.57

Active filters answer some of the limitations of the passive filter by offering isolation between stages and eliminating the need for

inductors. Their use at high frequencies is limited by the dynamic performance of the active elements.

ACTIVE FILTERS

- Eliminate Need for Inductors
- Good Interstage Isolation
- High Frequency Use Limited by Op Amp Dynamic Performance

Figure 2.58

A filter can be specified in terms of five parameters as shown in Figure 2.59. The *cutoff frequency* F_c is the frequency at which the filter response leaves the error band (or the -3dB point for a Butterworth filter). The *stopband frequency* F_s is the frequency at which the minimum attenuation in the

stopband is reached. The *passband ripple* A_{\max} is the variation (error band) in the passband response. The *minimum passband attenuation* A_{\min} defines the signal attenuation within the *stopband*. The *order M* of the filter is the number of poles in the transfer function.

KEY FILTER DESIGN PARAMETERS




Figure 2.59

FILTER SPECIFICATIONS

- Cutoff Frequency, F_c
- Stopband Frequency, F_s
- Passband Ripple, A_{\max}
- Stopband Attenuation, A_{\min}
- Filter Order, M

Figure 2.60

2

Typically, one or more of the above parameters will be variable. For instance, if you were to design an antialiasing filter for an ADC you will know the cutoff frequency, the stopband frequency, and the minimum attenuation. You can then go to a chart or computer program to determine the other parameters.

There are many transfer functions that may satisfy the requirements of a particular filter. The *Butterworth* filter is the best compromise between attenuation and phase response. It has no ripples in the passband or the stopband and is called the *maximally flat filter* because of this. The Butterworth filter achieves its flatness at the expense of a relatively wide transition region from passband to stopband.

The *Chebyshev* filter has a smaller transition region than the same-order Butterworth

filter, but it has ripples in either its passband or stopband. This filter gets its name because the Chebyshev filter minimizes the height of the maximum ripple—this is the Chebyshev criterion.

The Butterworth filter and the Chebyshev filter are all-pole designs. By this we mean that the zeros of the transfer function are at one of the two extremes of the frequency range (0 or ∞). For a lowpass filter the zeros are at $f = \infty$. We can add finite frequency transfer function zeros as well as poles to get an *Elliptical Filter*. This filter has a shorter transition region than the Chebyshev filter because it allows ripple in both the stopband and passband. The Elliptical filter also has degraded phase (time domain) response.

These are by no means all possible transfer functions, but they do represent the most common.

POPULAR FILTER DESIGNS

- **Butterworth:** All Pole, No Ripples in Passband or Stopband, Maximally Flat Response
- **Chebyshev:** All Pole, Ripple in Passband, Shorter Transition Region than Butterworth for Given Number of Poles
- **Elliptical:** Ripple in Both Passband and Stopband, Shorter Transition Region than Chebyshev, Degraded Phase Response, Poles and Zeros

Figure 2.61

Once the order of the filter and the specifications of filter have been determined, the design charts (see Reference 10) or computer programs are consulted, and the linear and quadratic factors of poles for the transfer function are determined. All filters, regardless of order, are made up of one- or two-pole sections. The single pole section is defined by its resonant frequency, which is the -3dB point. The pole pair in a two-pole filter section is defined by its resonant frequency (F_o) and Q , which indicates the peaking of the section. Sometimes alpha (α) is used instead of Q ($Q=1/\alpha$).


Armed with the various values F_o and Q , you then choose the configuration for the realization of the filter: Butterworth, Chebyshev, or Elliptical.

For passive filters, these values, along with the filter characteristic impedance determine the inductor, capacitor, and resistor values.

For active filters, you must decide which of the realizations you are going to use. The three most common are the *Sallen-Key* (voltage controlled voltage source), *multiple feedback*, and *state variable*. Each realization has its own advantages and disadvantages.

The Sallen-Key configuration shown in Figure 2.62 is the least dependent on the performance of the op amp, and the signal phase is maintained. For this filter the ratio of the largest resistor value to the smallest resistor value and the ratio of the largest capacitor value to the smallest capacitor value is low. The frequency term and Q terms are somewhat independent, but they are very sensitive to the gain parameter. The Sallen-Key is very Q -sensitive to element values for high Q sections. The design equations are also given in Figure 2.62.

VOLTAGE CONTROLLED VOLTAGE SOURCE (SALLEN-KEY) REALIZATION



H = Circuit Gain Below Cutoff

α = Damping Ratio = $1/Q$

F_o = Cutoff Frequency

Choose $C1$

$$K = 2\pi F_o C1$$

$$M = \frac{\alpha^2}{4} + H - 1$$

$$C2 = M C1$$

$$R1 = \frac{2}{K\alpha}$$

$$R2 = \frac{\alpha}{2MK}$$

Choose $R5$

$$R4 = R5(H - 1)$$

For $H = 1$, $R4 = 0$, $R5 = \text{Open}$

Figure 2.62

The multiple feedback realization shown in Figure 2.63 uses an op amp in the inverting configuration. The dependence on the op amp parameters are greater than in the Sallen-Key realization. It is hard to generate high Q sections due to the limitations of the open loop gain of the op amp. The maximum to minimum component value ratios are higher than in the Sallen-Key realization. The design equations are also given in Figure 2.63.

The state-variable realization shown in Figure 2.64 offers the most precise implementation, at the expense of many more circuit elements. All parameters can be adjusted independently, and lowpass, highpass, and bandpass outputs are all available simultaneously. The gain of the filter is also independently variable. Since all parameters of the state variable filter can be adjusted independently, component spread is minimized. Also variations due to temperature and component tolerances are minimized. The design equations for the state variable filter are given in Figure 2.64.

Another active filter realization that has recently become more popular is the *Frequency Dependent Negative Resistor* (FDNR), which is a subset of the *General Impedance Converter* (GIC). In the FDNR the passive realization goes through a transformation by $1/s$. Therefore inductors, whose impedance is sL , transform into a resistor of value L . Similarly, a resistor of value R becomes a capacitor of value R/s . A capacitor of impedance $1/sC$ transforms into a frequency dependent variable resistor, which is given the

designation D . Its impedance is $1/s^2C$. The transformations to the FDNR configuration and the GIC implementation of the D element are given in Figure 2.65.


The advantage of the FDNR realization is that there are no op amps in the signal path which can add noise. This realization is also relatively insensitive to component variation. The advantages of the FDNR come at the expense of an increase in the number of components required.

For all of the realizations discussed above, the tabulated filter values are in terms of the lowpass function normalized to a frequency of 1 radian/second with an impedance level of 1. To realize the final design, the filter values are scaled by the appropriate frequency and impedance.

Similarly, the lowpass prototype is converted to a highpass filter by scaling by $1/s$ in the transfer function. In practice this amounts to capacitors becoming inductors with a value $1/C$ and inductors becoming capacitors with a value of $1/L$ for passive designs. For active designs resistors become capacitors with a value of $1/R$, and capacitors become resistors with a value of $1/C$.

Transformation to the bandpass response is a little more complicated. If the corner frequencies of the bandpass are widely separated (by more than 2 octaves) the filter is made up of separate lowpass and highpass sections. In the case of a narrowband bandpass filter the design is much more complicated and is usually done using a computer program or design tables.

MULTIPLE FEEDBACK REALIZATION



F_O = Cutoff Frequency

α = Damping Ratio = $1/Q$

H = Absolute Value of Circuit Gain

Choose $C5$

$$K = 2\pi F_O C1$$

$$C2 = \frac{4C5}{\alpha^2} (H + 1)$$

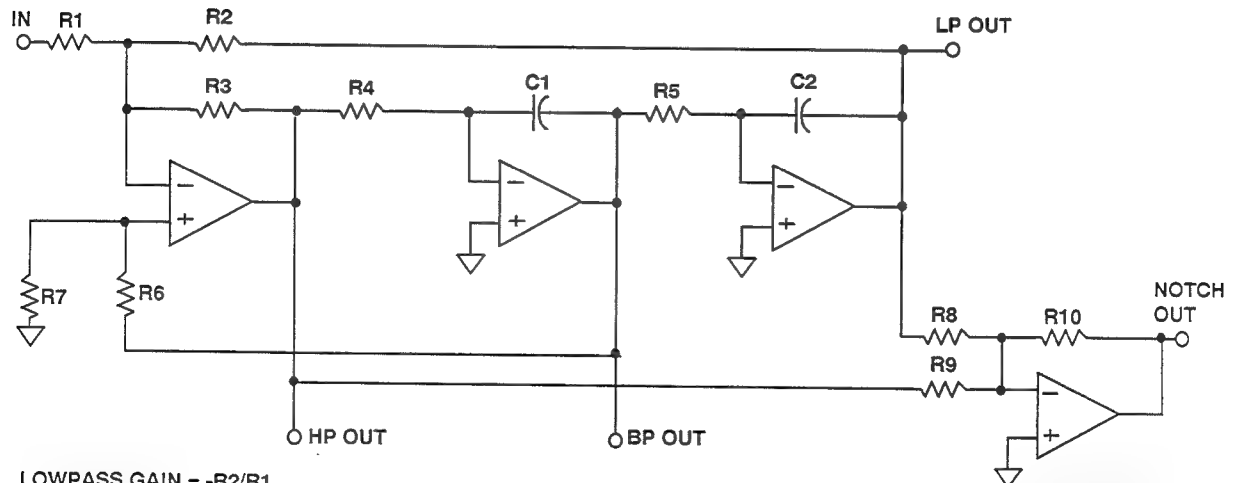
$$R1 = \frac{\alpha}{2HK}$$

$$R3 = \frac{\alpha}{2K(H + 1)}$$

$$R4 = HR1$$

Figure 2.63

STATE VARIABLE REALIZATION



LOWPASS GAIN = $-R2/R1$

HIGHPASS GAIN = $-R3/R1$

$$\text{BANDPASS GAIN} = \frac{R6 + R7}{R1R7 \left(\frac{1}{R1} + \frac{1}{R2} + \frac{1}{R3} \right)}$$

$$F_O = \frac{1}{2\pi} \sqrt{\frac{R3}{R2 \cdot R4 \cdot R5 \cdot C1 \cdot C2}}$$

$$Q = \frac{1}{\alpha} = \frac{R6 + R7}{R7} \left(\frac{1}{\frac{1}{R1} + \frac{1}{R2} + \frac{1}{R3}} \right) \sqrt{\frac{R4 \cdot C1}{R2 \cdot R3 \cdot R5 \cdot C2}}$$

FOR NOTCH FREQUENCY = F_Z

$$\text{FOR } F_O = F_Z, \frac{R2 \cdot R9}{R3 \cdot R8} = 1$$

$$\text{FOR } F_O > F_Z, \frac{R2 \cdot R9}{R3 \cdot R8} < 1$$

$$\text{FOR } F_O < F_Z, \frac{R2 \cdot R9}{R3 \cdot R8} > 1$$

$$\frac{F_Z^2}{F_O^2} = \frac{R2 \cdot R9}{R3 \cdot R8}$$

Figure 2.64

FREQUENCY DEPENDENT NEGATIVE RESISTOR 1/S IMPEDANCE TRANSFORMATION




Figure 2.65

SOME ACTIVE FILTER REALIZATIONS

- **Sallen-Key:** Good Phase Response, Least Dependent on Op Amp Performance, Sensitive to Element Values for High Q Sections
- **Multiple Feedback:** Less Sensitive to Element Values, High Q Sections Difficult due to Op Amp Open Loop Gain Limitations
- **State-Variable:** Most Precise, More Components, All Parameters Independently Adjustable
- **Frequency Dependent Negative Resistance (FDNR):** Op Amps not in Signal Path, More Components, Relatively Insensitive to Component Variations

Figure 2.66

ANTI_ALIASING FILTER DESIGN EXAMPLE

We will now design a passive and active antialiasing filter based upon the same specifications. The active filter will be designed in four realizations: Sallen-Key, multiple feedback, state variable, and Frequency Dependent Negative Resistance (FDNR). We

choose the Butterworth filter in order to give the best compromise between attenuation and phase response.

The specifications for the filter are as follows:

ANTI_ALIASING FILTER SPECIFICATIONS

- Cutoff Frequency $F_c = 8\text{kHz}$
- Stopband Attenuation F_s at $50\text{kHz} = 70\text{dB}$
- Best Balance Between Attenuation and Phase Response
- Choose Butterworth Design
- From Design Charts, for $f = 6.25$ ($50\text{kHz}/8\text{kHz}$), $M = 5$

Figure 2.67

Consulting the design charts (Reference 10, p. 82), we see that for 70dB of attenuation at a frequency of 6.25 ($50\text{kHz}/8\text{kHz}$) a

fifth order filter is required.

We now consult the tuning tables (Reference 10, p. 341) and find:

ALPHA AND F_0 VALUES FROM TUNING TABLES

STAGE	ALPHA	F_0
1	---	1.000
2	1.618	1.000
3	0.618	1.000

Figure 2.68

The first stage is a real pole, thus the lack of an alpha value. It should be noted that this is not necessarily the order of implementation in hardware. In general you would typically put the real pole last and put the second order sections in order of decreasing alpha (increasing Q).

For the passive design we will choose the zero input impedance configuration. From the design table (Reference 10, p. 313) we find the following normalized values for the filter:

NORMALIZED PASSIVE FILTER VALUES FROM TABLES

$$\begin{aligned} L1 &= 1.5451 & C2 &= 1.6944 \\ L3 &= 1.3820 & C4 &= 0.8944 \\ L5 &= 0.3090 \end{aligned}$$

Figure 2.69

These values are for a 1 rad/second filter with a 1 ohm termination. To scale the filter we divide all reactive elements by the desired cutoff frequency, 8kHz (50265 rad/sec). We also need to scale the impedance. For this example, we choose a value of 1000 ohms. To scale the impedance we multiply all resistor and inductor values and divide all capacitor values by the impedance scaling factor. After scaling, the circuit looks like Figure 2.70.

For the Sallen-Key active realization, we use the design table shown in Figure 2.62. The values for C1 in each section are chosen to give reasonable resistor values. The implementation is shown in Figure 2.71. For the Sallen-Key realization to work correctly, it is assumed to have a zero-impedance driver and a return path for dc. Both of these criteria are approximately met when you use an op amp to drive the filter.

EXAMPLE FILTER PASSIVE IMPLEMENTATION




Figure 2.70

EXAMPLE FILTER SALLEN-KEY IMPLEMENTATION

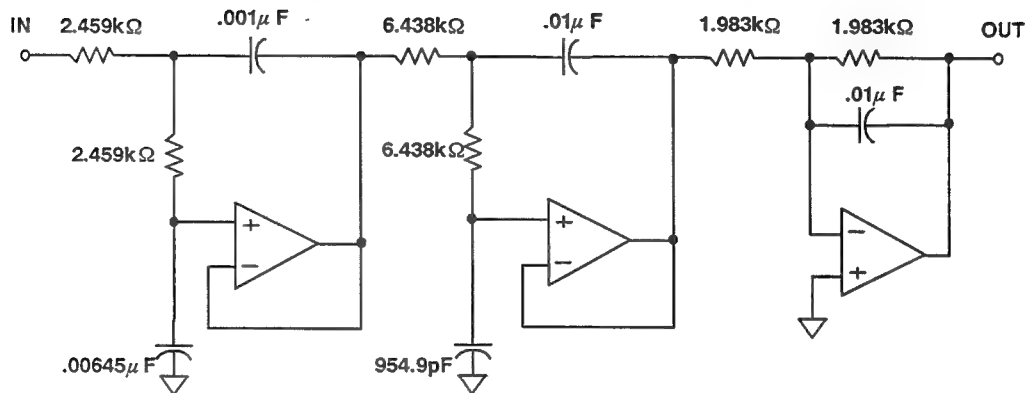


Figure 2.71

EXAMPLE FILTER MULTIPLE FEEDBACK IMPLEMENTATION




Figure 2.72

EXAMPLE FILTER STATE VARIABLE IMPLEMENTATION

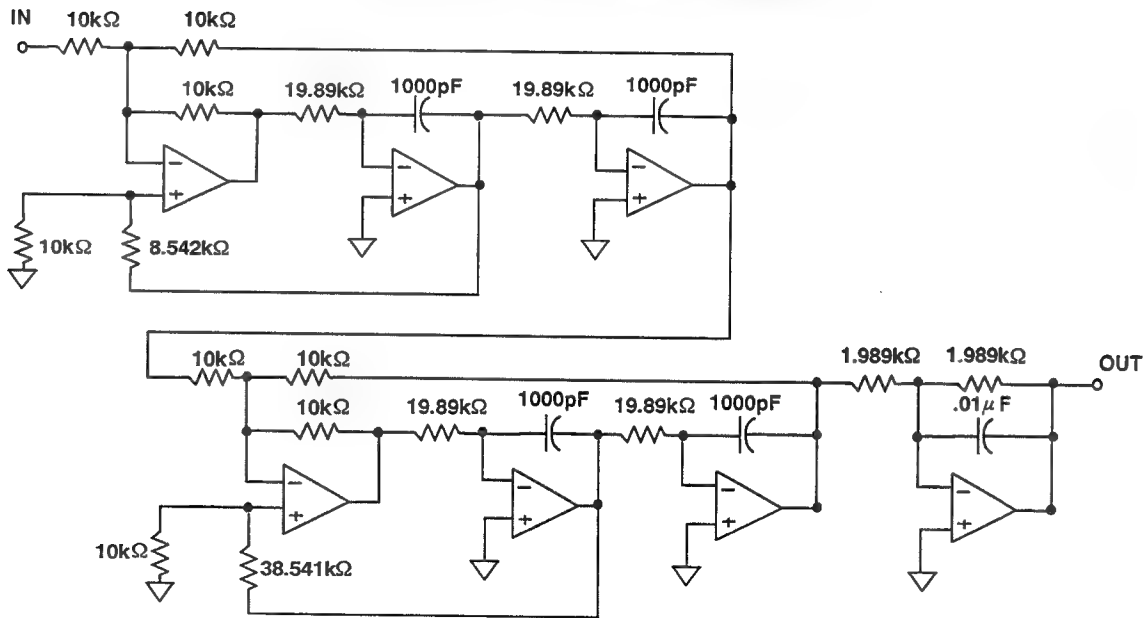


Figure 2.73

EXAMPLE FILTER FDNR IMPLEMENTATION

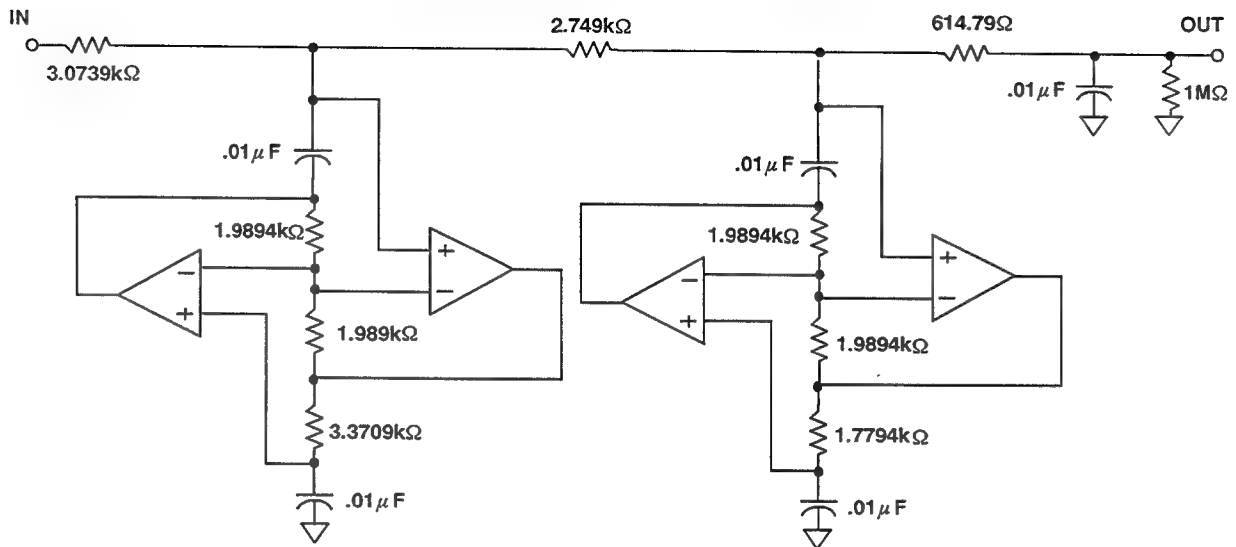


Figure 2.74

Figure 2.72 shows a multiple feedback realization of our filter. It was designed using the equations in Figure 2.63.

The state variable realization is shown in Figure 2.73, and the Frequency Dependent Negative Resistance (FDNR) realization is shown in Figure 2.74. In the conversion process from passive to FDNR, the D element is normalized for a capacitance of 1F. We then scale the filter to a more reasonable value (0.01μF in this case).

In all of the filters above the values shown are the exact calculated values. These exact values are rarely obtainable. We must therefore either substitute the nearest standard value or use series/parallel combinations. Any variation from the ideal values will cause a shift in the filter response characteristic, but often the effects are minimal.

A PROGRAMMABLE STATE VARIABLE FILTER

A realization of a programmable state variable filter using DACs is shown in Figure 2.75. DACs A1 and B1 control the gain and Q of the filter characteristic, while DACs A2 and B2 must accurately track for the simple expression for f_c to be true. This is readily accomplished using two AD7528 DACs and one AD713 quad op amp. Capacitor C3

The computer can be used to evaluate these variations on the overall performance and determine if they are acceptable.

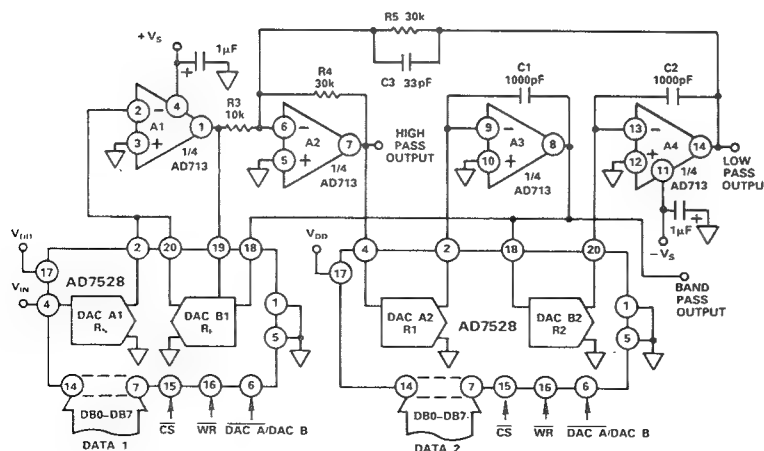
In active filter applications using op amps, the dc accuracy of the amplifier is often critical to optimal filter performance. The amplifier's offset voltage will be passed by the filter and may be amplified to produce excessive output offset. For low frequency applications requiring large value resistors, bias currents flowing through these resistors will also generate an output offset voltage.

In addition, at higher frequencies, an op amp's dynamics must be carefully considered. Here, slewrate, bandwidth, and open loop gain play a major role in op amp selection. The slewrate must be fast as well as symmetrical to minimize distortion.

compensates for the effects of op amp and gain-bandwidth limitations.

This filter provides lowpass, highpass, and bandpass outputs and is ideally suited for applications where digital control of filter parameters is required. The programmable range for component values shown is $f_c = 0$ to 15kHz, and $Q = 0.3$ to 4.5.

A PROGRAMMABLE STATE VARIABLE FILTER CIRCUIT



CIRCUIT EQUATIONS

$$C_1 = C_2, R_1 = R_2, R_4 = R_5$$

$$f_c = \frac{1}{2\pi R_1 C_1}$$

$$Q = \frac{R_3}{R_4} \cdot \frac{R_F}{R_{FBB1}}$$

$$A_0 = -\frac{R_F}{R_S}$$

NOTE:

DAC equivalent resistance equals $\frac{256 \times (\text{DAC Ladder resistance})}{\text{DAC Digital Code}}$

Figure 2.75

SEVEN-POLE FDNR 20kHz ANTIALIASING FILTER

Figure 2.76 shows a 7-pole antialiasing filter for a 2x oversampling (88.2kSPS) digital audio application. This filter has less than 0.05dB passband ripple and $19.8 \pm$

0.3 μ s delay, dc-20kHz. The filter will handle a 5V rms signal ($V_s = \pm 15$ V) with no overload at any internal nodes. The frequency response of the filter is shown in Figure 2.77.

20kHz FDNR AUDIO ANTIALIASING FILTER

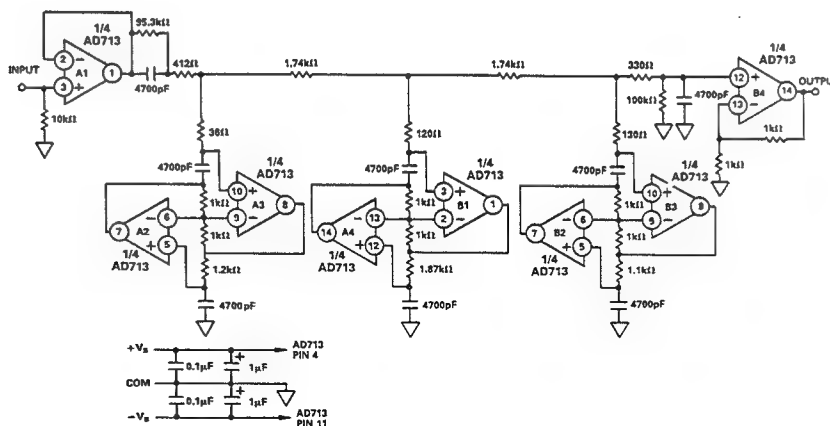


Figure 2.76

AUDIO ANTIALIASING FILTER RESPONSE




Figure 2.77

A WIDEBAND SALLEN-KEY FILTER

Figure 2.78 shows an AD843 FET input op amp used in a 1MHz Sallen-Key filter. This circuit also works well with the AD841, AD845, or AD847. The circuit is designed to

be a maximum-flatness filter with a Q of 0.575 and a dc gain of 1.26. The frequency response of the filter to a 0dBm input signal is shown in Figure 2.79.

1 MHz SALLEN KEY FILTER

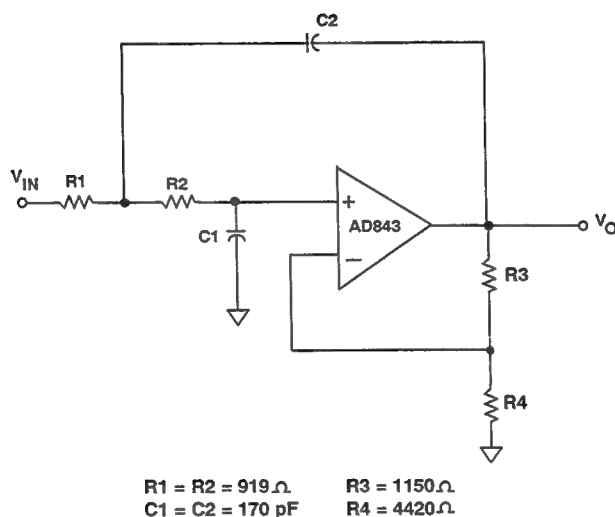


Figure 2.78

SALLEN-KEY SMALL SIGNAL FREQUENCY RESPONSE




Figure 2.79

REFERENCES

1. Robert M. Clarke, Industry's Fastest Monolithic FET-Input Op Amp, **Analog Dialogue** 23-4, 1989, p.13.
2. Barrie Gilbert, **ISSCC Digest of Technical Papers 1968**, pp. 114-115 February 16, 1968
3. Barrie Gilbert, **Journal of Solid State Circuits**, Vol. SC-3, Dec. 1968, pp. 353-372.
4. Daniel H. Sheingold, Editor, **Nonlinear Circuits Handbook**, Analog Devices, 1974.
5. Charles Kitchin and Lew Counts, **RMS to DC Conversion Application Guide**, Second Edition, Analog Devices, 1986.
6. Richard Smith Hughes, **Logarithmic Amplifiers**, Artech House, Inc., Dedham, MA., 1986
7. William L Barber and Edmund R. Brown, A True Logarithmic Amplifier for Radar IF Applications, **IEEE Journal of Solid State Circuits**, Vol. SC-15, No. 3, June 1980, pp. 291-295.
8. **Broadband Amplifier Applications**, Plessey Co. Publication P.S. 1938, Sept. 1984.
9. **High Speed Design Seminar**, Analog Devices, 1990.
10. A. I. Zverev, **Handbook of Filter Synthesis**, John Wiley, 1967.
11. **AD7528 Dual 8-Bit CMOS DAC Application Note**, Analog Devices.
12. **8th Order Programmable Lowpass Analog Filter Using 12-Bit DACs**, Analog Devices.
13. **CMOS DAC Application Guide**, Analog Devices.
14. A. B. Williams, **Electronic Filter Design Handbook**, McGraw-Hill, 1981.
15. M. E. Van Valkenburg, **Introduction to Modern Network Synthesis**, John Wiley and Sons, 1960.
16. A. I. Zverev and H. J. Blinichikoff, **Filtering in the Time and Frequency Domain**, John Wiley and Sons, 1976.
17. A. Budak, **Passive and Active Network Analysis and Synthesis**, Houghton Mifflin Company, Boston, 1974.

18. L. P. Huelsman and P. E. Allen, **Introduction to the Theory and Design of Active Filters**, McGraw Hill, 1980.
19. M. S. Ghausi and K. R. Laker, **Modern Filter Design: Active RC and Switched Capacitors**, Prentice Hall, 1981.

SECTION III

3

FUNDAMENTALS OF SAMPLED DATA SYSTEMS

FUNDAMENTALS OF SAMPLED DATA SYSTEMS

- A TYPICAL DSP SAMPLED DATA SYSTEM
- DISCRETE TIME SAMPLING OF ANALOG SIGNALS
- SELECTION OF ANTIALIASING FILTER
- OVERSAMPLING AND DECIMATION
- UNDERSAMPLING AND ITS APPLICATIONS
- EFFECTS OF FINITE AMPLITUDE RESOLUTION DUE TO QUANTIZATION
- QUANTIZATION THEORY, SIGNAL TO NOISE RATIO, AND EFFECTIVE BITS
- SELECTION OF ADC RESOLUTION BASED ON SIGNAL DYNAMIC RANGE
- ADC STATIC TRANSFER CHARACTERISTICS
- DAC STATIC TRANSFER CHARACTERISTICS
- ADC DYNAMIC PERFORMANCE

Signal to Noise Ratio and Effective Bits

Peak Spurious, Peak Harmonic Content, and Spurious Free Dynamic Range (SFDR)

Total Harmonic Distortion

Full Power Bandwidth

Full-Linear Bandwidth

Intermodulation Distortion (IMD)

AC Linearity Plots Using Histograms

Aperture Delay Time (or Effective Aperture Delay Time)

Aperture Jitter

Transient Response or Settling Time

Overvoltage Recovery

■ **DAC DYNAMIC PERFORMANCE**

Settling Time

Glitch Impulse Area

Harmonic Distortion

Deglitching DACs Using SHAs

$\sin(x)/x$ Frequency Rolloff Effect

■ **SWITCHED CAPACITOR FILTERS**

SECTION III

FUNDAMENTALS OF SAMPLED DATA SYSTEMS

A TYPICAL DSP SAMPLED DATA SYSTEM

A block diagram of a typical sampled data DSP system is shown in Figure 3.1. Prior to the actual analog-to-digital conversion, the analog signal usually passes through some sort of signal conditioning circuitry which performs such functions as amplification, attenuation, or filtering. If the analog signal originates as a temperature, pressure, flow-rate, or force, then an appropriate sensor and

transducer is required to first convert the physical quantity into an electrical voltage or current.

There are two key concepts involved in the actual analog-to digital conversion process: *discrete time sampling* and *finite amplitude resolution due to quantization*. An understanding of these concepts is vital to DSP applications.

KEY ELEMENTS OF A SAMPLED DATA SYSTEM

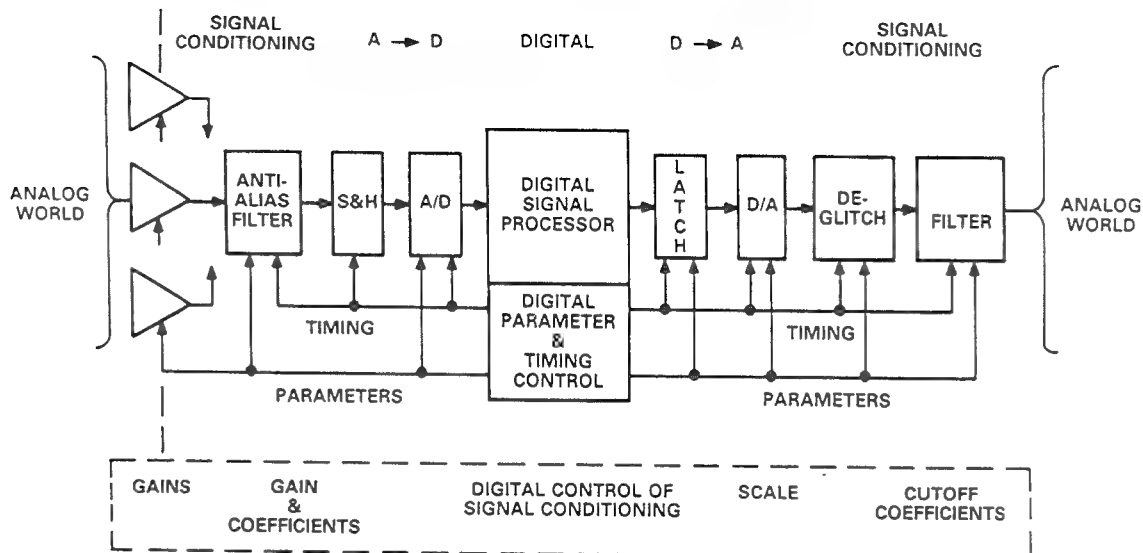


Figure 3.1

DISCRETE TIME SAMPLING OF ANALOG SIGNALS

The concept of discrete time and amplitude sampling of an analog signal is shown in Figure 3.2. The continuous analog data must be sampled at discrete intervals, t_s , which must be carefully chosen to insure an accurate representation of the original analog signal. It is clear that the more samples

taken (faster sampling rates), the more accurate the digital representation, but if fewer samples are taken (lower sampling rates), a point is reached where critical information about the signal is actually lost. This leads us to the statement of Nyquist's criteria given in Figure 3.3.

DISCRETE SAMPLING OF AN ANALOG SIGNAL




Figure 3.2

NYQUIST'S CRITERIA

- An Analog Signal with a *Bandwidth* of f_a Must be Sampled at a Rate $f_s > 2f_a$ in Order to Avoid the Loss of Information
- If $f_s < 2f_a$, then a Phenomena Called Aliasing Will Occur in the Analog Signal Bandwidth

Figure 3.3

In order to understand the implications of *aliasing* in both the time and frequency domain, first consider the four cases of a time domain representation of a sampled sine-wave signal shown in Figure 3.4. In the Case 1, it is clear that an adequate number of samples have been taken to preserve the information about the sinewave. In Case 2 of the figure, only four samples per cycle are taken; still an adequate number to preserve the information. Case 3 represents the ambiguous limiting condition where $f_s = 2f_a$. If the relationship between the sampling points and the sinewave were such that the sine-wave was being sampled at precisely the zero crossings (rather than at the peaks, as shown in the illustration), then all information regarding the sinewave would be lost. Case 4 of Figure 3.4 represents the situation where $f_s < 2f_a$, and the information obtained from the samples indicates a sinewave having a frequency which is lower than $f_s/2$, i.e. the out-of-band signal is *aliased* into the Nyquist bandwidth between dc and $f_s/2$. As the sampling rate is further decreased, and the analog input frequency f_a approaches the sampling frequency f_s , the aliased signal approaches dc in the frequency spectrum.

The corresponding frequency domain representation of the above scenario is shown in Figure 3.5. Note that sampling the analog signal f_a at a sampling rate f_s actually pro-

duces two alias frequency components, one at $f_s + f_a$, and the other at $f_s - f_a$. The upper alias, $f_s + f_a$, seldom presents a problem, since it lies outside the Nyquist bandwidth. It is the lower alias component, $f_s - f_a$, which causes problems when the input signal exceeds the Nyquist bandwidth, $f_s/2$.

It is clear from the above discussion that the ADC must be preceded by an anti-aliasing filter which has sufficient stopband attenuation at $f_s/2$ and above to prevent unwanted in-band aliasing. Aliasing may also occur from harmonics of the fundamental signal which fall outside the Nyquist bandwidth, or from unfiltered broadband noise at the ADC input.

The effects of aliasing on the dynamic range of a sampled data system are shown in Figure 3.6. The top part of the figure illustrates the desired condition at the Nyquist point, where the aliased component intersects the input signal at a point below the desired dynamic range. The lower part of the figure shows the condition where the upper-frequency dynamic range is limited by the aliased components. This condition will result in a reduction in overall signal-to-noise ratio at the higher frequencies, and could result in the distortion due to aliased out-of-band tones or harmonics as shown in Figure 3.7.

TIME DOMAIN EFFECTS OF ALIASING




Figure 3.4

FREQUENCY DOMAIN EFFECTS OF ALIASING

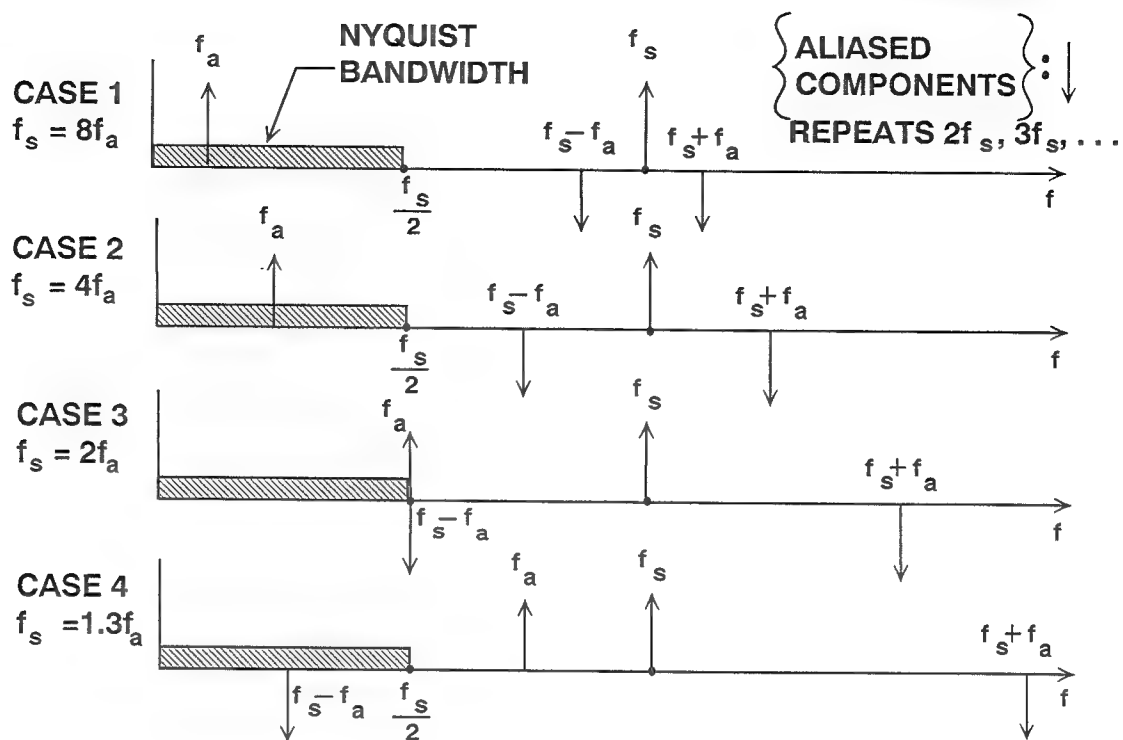


Figure 3.5

FREQUENCY DOMAIN EFFECTS OF ALIASING ON DYNAMIC RANGE




Figure 3.6

UNWANTED TONES DUE TO ALIASING




Figure 3.7

SELECTION OF ANTIALIASING FILTERS

It should be clear by now, that for a given analog input bandwidth, f_a , the requirements of the antialiasing filter are related not only to the sampling rate, f_s , but also to the desired system dynamic range. Simply stated, *dynamic range* is the ratio of the largest expected signal to the smallest signal which must be resolved, and is usually expressed in dB. At this point, we are concerned with dynamic range limitations due to aliasing. The limiting effects of ADC quantization noise and other non-linearities will be discussed shortly. The following rules of thumb will result in a filter which is somewhat overspecified, but the concepts are valid and can be refined to fit the actual system requirements.

First, set the corner frequency of the antialiasing filter equal to the desired analog input bandwidth, f_a . This defines the pass-

band of the filter, $f_{\text{pass}} = f_a$. Define the beginning of the filter's stopband, $f_{\text{stop}} = f_s/2$. Let the filter stopband attenuation be the desired upper-frequency dynamic range, DR, expressed in dB. These parameters define the transition band characteristics of the filter, i.e., it must achieve a stopband attenuation equal to the dynamic range over $\log_2(f_{\text{stop}}/f_{\text{pass}})$ octaves. The approximate order of the filter, M , (the number of poles) required to achieve this transition band slope can then be determined, since the filter rolloff is approximately 6M dB per octave. A simple example calculation is shown in Figure 3.9, where the signal bandwidth, f_a , is 3kHz, the sampling rate, f_s , is 12kHz, and a dynamic range of 60dB is required. This implies that a 10 pole filter is needed. Remember that in practice, any analog filter with more than 8 poles becomes a real design

ANTIALIASING FILTER REQUIREMENTS




Figure 3.8

ANTIALIASING FILTER EXAMPLE




Figure 3.9

EFFECTS OF OUT-OF-BAND ATTENUATION ON ANTIALIASING FILTER




Figure 3.10

challenge, and a filter with more than 12 poles becomes almost an impossibility except for the experienced filter designer. These considerations so far have neglected the filter's phase characteristics, and also the in-band and out-of-band ripple requirements. The addition of these parameters can make antialiasing filter design a truly formidable task.

The above rules-of-thumb for determining the complexity of the antialiasing filter assume that fullscale signals can occur at essentially all input frequencies above

Nyquist. In actual practice, this is not usually the case, and there is some natural attenuation of the signal being processed at the higher input frequencies. For instance,

in the previous example, if signals at Nyquist and above were already attenuated by 12dB, then a filter stopband attenuation of only 48dB would be required at the Nyquist frequency of 6kHz. This would imply that only an 8 pole filter would be needed. This situation is illustrated in Figure 3.10.

From the above discussions, it is clear that the requirements on the antialiasing filter can be relaxed at the expense of higher sampling rates (called *oversampling*). Later in the seminar, we will see that a particular class of ADCs and DACs, called *Sigma-Delta* ($\Sigma\Delta$) are inherently oversampling converters and greatly reduce the complexity of the antialiasing filter.

OVERSAMPLING AND DECIMATION

As previously discussed, the major advantage of oversampling the input signal is the resulting simplification in the antialiasing filter requirements. Of course the downside of oversampling is that it also increases the ADC output data rate, and the DSP must be able to keep up in order to maintain real-time operation. If the data is to be transmitted in serial form, then it consequently will occupy more of the frequency spectrum. An attractive alternative makes use of both analog *and* digital filtering techniques, oversampling, and a process called *decimation*. Figure 3.11 shows the traditional case, where all the antialiasing burden lies with the analog input filter preceding the ADC. In Figure 3.12, however, the oversampling ratio, K (K is an integer), relaxes the rolloff requirement of the input analog filter by increasing the Nyquist frequency to $Kf_s/2$. The digital filter following the ADC (digital filtering will be discussed at length in Section VII) implements the antialiasing function with respect to f_s , and has sufficient stopband attenuation at $f_s/2$ to achieve the desired

dynamic range. As we will see later in the seminar, digital filters having sharp cutoff characteristics with good phase response are much more easily implemented than their corresponding analog counterparts (assuming sufficient speed in the DSP). Finite Impulse Response (FIR) filters can be designed which have linear phase characteristics. Since the bandwidth has been reduced to $f_s/2$ by the digital antialiasing filter, the data coming out of the digital filter actually contains redundant information, and there is no need to look at every sample. In fact, it is only necessary to look at every K th sample. This process is called *decimation*, and will be discussed in much more detail in the section of the seminar on Sigma-Delta converters (Section VI). In addition, the actual decimation can be performed by the FIR filter itself by computing a single output sample for every K input samples. This concept of oversampling and decimation is one of the most powerful concepts in real-world DSP.

NYQUIST SAMPLING WITH ANALOG LOWPASS FILTER

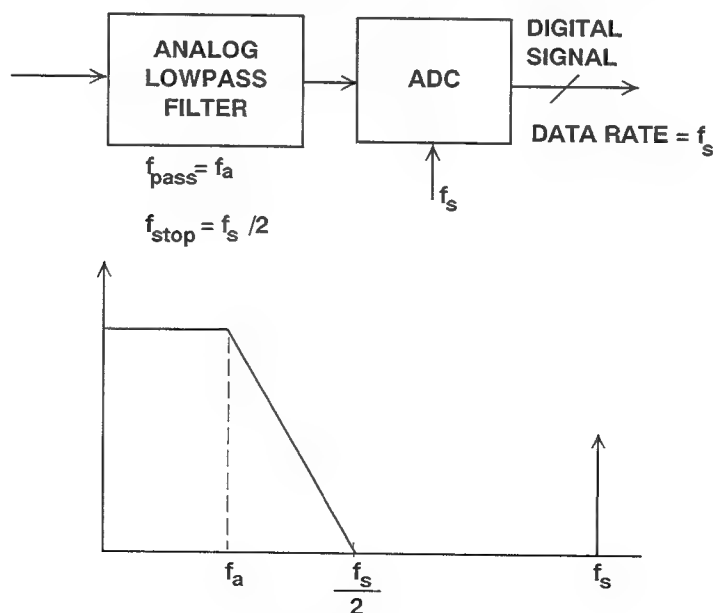


Figure 3.11

OVERSAMPLING WITH ANALOG AND DIGITAL FILTERING

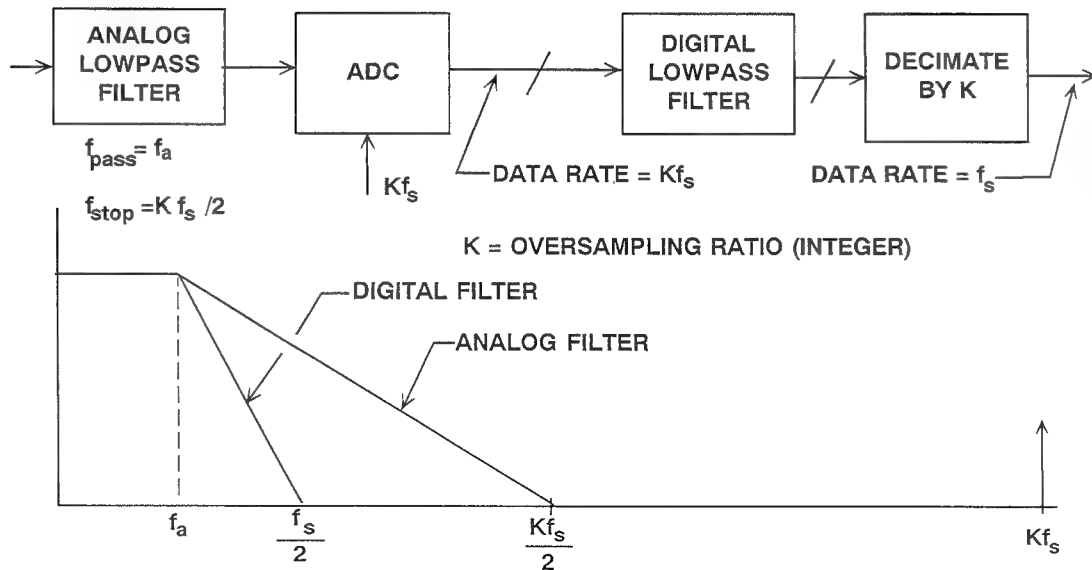


Figure 3.12

UNDERSAMPLING AND ITS APPLICATIONS

In this section we will see that there are some applications in DSP where aliasing is perfectly acceptable and can be used advantageously. When the analog signal being digitized by the ADC exceeds $f_s/2$, the condition is often referred to as *super-Nyquist*, or *undersampling*. Nyquist's criteria states that the *bandwidth* (not the actual frequency) of the signal being digitized should not exceed $f_s/2$ for information to be preserved. As an example, consider a telecommunications transmultiplexer application where Frequency Division Multiplexed (FDM) data occupying the bandwidth of 60 to 108kHz is sampled at a frequency of 112kHz. Figure 3.13 shows the spectrum of the signal and the location of the aliased components. At the receiving end of the system, the filter which follows the reconstruction DAC is a bandpass rather than a lowpass and must

filter out the aliased components falling between 4kHz and 52kHz as well as the component located at the sampling frequency of 112kHz.

Another application for super-Nyquist operation is in the direct conversion of IF signals to baseband. Most traditional communication and radar receivers employing ADCs and DSP utilize a system in which the intermediate frequency (IF) from the front end of the receiver is down-converted or demodulated to a baseband signal by a mixer and a lowpass filter as shown in Figure 3.14. This final IF stage uses a local oscillator which is phase coherent with the signal carrier frequency. The mixer output contains a baseband signal which is proportional to the phase difference between the two inputs. Following the mixer is a lowpass filter, amplifier, and an ADC. Typical mixers have

a conversion loss ranging from 4 to 6dB. In cases when the signal-to-noise ratio is limited by the front end, elimination of the mixer

will improve the overall noise figure of the receiver.




Figure 3.13

ANALOG DOWNCONVERSION OR DEMODULATION

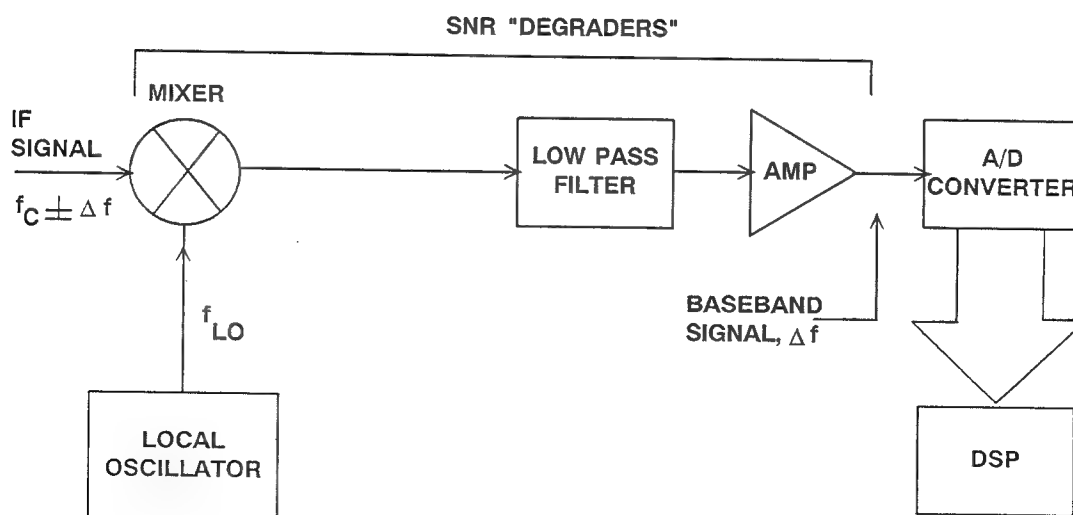


Figure 3.14

DIRECT IF TO DIGITAL DOWN CONVERSION OR DEMODULATION

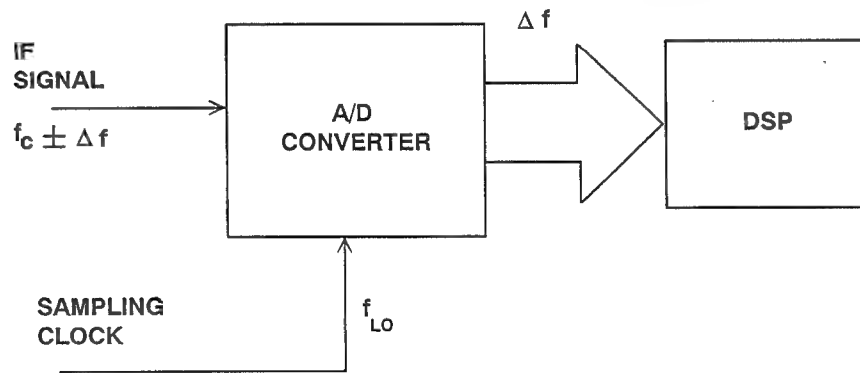


Figure 3.15

This can be accomplished (as shown in Figure 3.15) if the IF frequency is sampled at a rate equal to the local oscillator frequency. The ADC now functions as a demodulator. If the ADC samples an analog signal of the same frequency as the sampling frequency, the digitized output is a dc value. Any deviation in the analog signal from the sampling frequency looks like a *beat* frequency, Δf , and the demodulation process is thereby achieved.

The data from the ADC must be processed by the DSP using an FFT which computes both the real and imaginary components of the digitized signal. This is necessary in order to preserve the phase information

contained in the demodulated signal.

Operation of ADCs in a super-Nyquist environment obviously requires that the dynamic performance of the converter be known for input frequencies *above* Nyquist. The signal-to-noise ratio and harmonic distortion performance of an ADC typically degrades at higher input frequencies, so ac performance for the input frequency desired must be adequate to meet system requirements. Super-Nyquist operation typically requires an ADC which is more robust to high frequency input signals than an ADC which is specified for strictly sub-Nyquist applications.

EFFECTS OF FINITE AMPLITUDE RESOLUTION DUE TO QUANTIZATION

The second major effect to be considered in a sampled data system is that of the finite amplitude resolution caused by the analog-to-digital or digital-to-analog conversion process. In this discussion, we will refer to the number of bits of the ADC (or DAC) as the converter's *resolution*, N . In the case of an ADC, the input range is divided into 2^N discrete levels, each represented by an N -bit binary word. For a DAC, the input consists of an N -bit binary word, and there are 2^N possible discrete output levels. Figure 3.16 shows the number of bits, N , the corresponding number of levels, 2^N , and the weight of

the *least significant bit* (LSB) expressed as a percentage and a ratio in dB [$20 \log_{10}(2^N)$], or $6.02N$ dB. This ratio (whether expressed as a percentage or in dB) represents the *dynamic range* of the converter, i.e., the ratio of the largest resolvable signal to the smallest resolvable signal. At this point, we should point out that the dynamic range values in Figure 3.16 represent *ideal* ADCs and DACs does not consider such ac performance limitations such as harmonic and intermodulation distortion. Neither do these values represent the theoretical signal-to-quantization noise. These topics will be discussed shortly.

RESOLUTION AND DYNAMIC RANGE OF ADCs AND DACs

#BITS, N	#LEVELS, 2^N	%, $100/2^N$	dB, $6N$
8	256	0.4	48
10	1024	0.1	60
12	4096	0.025	72
14	16384	0.006	84
16	65536	0.0015	96
18	262144	0.0004	108
20	1048576	0.0001	120
22	4194304	0.000025	132
24	16777216	0.000006	144

Figure 3.16

QUANTIZATION THEORY, SIGNAL TO NOISE RATIO, AND EFFECTIVE BITS

The finite resolution of ADCs and DACs gives rise to a theoretical limitation to the signal-to-noise ratio (SNR) which is a function of the number of bits, N . In order to make a meaningful measurement, the ADC is stimulated with a fullscale sinewave input which is slightly below the clipping range of the converter. This gives rise to a sample-to-sample error which produces *quantization noise*. It can be shown mathematically that the rms noise voltage produced by quantization measured *within the Nyquist bandwidth*

is given by the familiar expression $q/\sqrt{12}$, where q is the weight of the least significant bit (LSB) of the converter. The value for the LSB, q , can be calculated by dividing the fullscale range of the ADC or DAC by 2^N . In an ideal converter with no error sources, the theoretical rms quantization noise voltage is also independent of both the input signal amplitude and frequency. The derivation for this simple expression is given in the following reference:

W.R. Bennett, *Spectra of Quantized Signals*, BSTJ 27, pp. 446-472, July 1948

For a fullscale sinewave input, it can further be shown that the theoretical rms signal to quantization noise ratio is given by $SNR = 6.02N + 1.76dB$.

QUANTIZATION THEORY BASICS

- RMS Quantization Noise in Nyquist Bandwidth, $f_s/2$:

$$q/\sqrt{12}$$

- Fullscale Sinewave RMS Signal to RMS noise ratio in Nyquist Bandwidth:

$$SNR = 6.02N + 1.76dB$$

- Effective Number of Bits (ENOB):

$$ENOB = \frac{SNR_{ACTUAL} - 1.76dB}{6.02}$$

Figure 3.17

It should be noted that the rms quantization noise generally approximates broadband noise across the Nyquist bandwidth. There are certain conditions, however, where this is not true. If there is correlation between the quantization error signal and the signal being digitized, then the quantization noise may be concentrated at harmonics of the input signal rather than being spread uniformly across the bandwidth. This is most likely to occur if the input signal is a sine-wave which is a subharmonic of the sampling frequency.

In testing ADCs, the SNR is usually calculated using DSP techniques while applying a pure sinewave signal to the input of the ADC as shown in Figure 3.18. The Fast Fourier Transform (FFT) processes a finite number of time samples and converts them into the frequency spectrum such as that shown in Figure 3.19 for the AD678 12-bit 200kSPS sampling ADC. The frequency spectrum is then used to calculate the SNR as well as harmonics of the fundamental input signal, very similar to an analog spectrum analyzer. The rms value of the signal is first computed. Then the rms value of all other frequency components over the Nyquist

bandwidth (this includes not only noise but also distortion products) is computed. The ratio of these two quantities, expressed in dB is the SNR. Various error sources in the ADC cause the measured SNR to be less than the theoretical value, $6.02N + 1.67dB$. These errors occur due to integral and differential nonlinearities, missing codes, and internal ADC noise sources. In addition, the errors typically are a function of input slewrate and therefore increase as the input frequency gets higher. In calculating the rms value of the noise, it is customary to include harmonics of the fundamental signal. This is sometimes referred to as the signal-to-noise-plus-distortion, $S/(N+D)$, but is usually called simply SNR. A typical plot of $S/(N+D)$ for the AD678 sampling ADC (12 bit, 200kSPS) is shown in Figure 3.20.

Another way to interpret SNR is in terms of *effective number of bits*, or *ENOBs*. The effective-bit calculation is performed by solving the SNR equation for N , given the measured value of SNR. (See Figure 3.17). For instance, a perfect 12 bit ADC would have a theoretical SNR of 74dB, corresponding to 12 effective bits. A measured SNR of 68dB, however, would correspond to 11 effective

ADC DYNAMIC TESTING




Figure 3.18

2048 POINT FFT OUTPUT FOR AD678 12-BIT, 200 kSPS ADC




Figure 3.19

S / (N + D) AND EFFECTIVE BITS FOR AD678 12-BIT, 200 kSPS ADC




Figure 3.20

bits. This says that the performance of the actual 12 bit ADC is equivalent to that of a perfect 11 bit ADC. Figure 3.20 also shows the ENOB performance of the AD678 on the same graph as the SNR. Note that at low frequencies, the AD678 exceeds 11.4 effective bits.

Effective bits can also be measured using the *sinewave curvefit* method. In this method, a sinewave is applied to the ADC, and a number of samples are collected. Instead of performing an FFT on the time samples, the *best-fit* sinewave to fit the data points is calculated. The sinewave amplitude, offset, frequency, and phase are chosen

to minimize the rms error between the actual sinewave data points and the theoretical sinewave. Again, the theoretical rms error for a perfect ADC is $q/\sqrt{12}$. The rms error between the actual sinewave and the theoretical sinewave is computed, and the effective bits are calculated using the formula shown in Figure 3.21. The ENOB measurement using the sinewave curvefit method correlates well with that obtained using the SNR technique. If the SNR calculation is made with a signal which is less than full-scale, then a correction factor must be added as shown in order for the two methods to correlate.

CALCULATING ENOB USING SINEWAVE CURVE FITTING

- Q_A = Actual RMS Error from Best Fit Sinewave
- Q_T = Theoretical N-Bit RMS Error from Best Fit Sinewave

$$= q/\sqrt{12}$$
- $ENOB = N - \log_2 \left[\frac{Q_A}{Q_T} \right]$ Correlates to:
- $SNR = \frac{SNR_{Actual} - 1.76dB + \text{Level of Signal Below FS}}{6.02}$

Figure 3.21

SELECTION OF ADC RESOLUTION BASED ON SIGNAL DYNAMIC RANGE

Selection of the proper ADC for a given application involves much more than just determining the number of bits required and the sampling rate. The dc and ac characteristics of the ADC must be examined with respect to

the analog signal being processed and a proper match must be found. Inevitably, this process involves certain tradeoffs in performance and cost.

DSP APPLICATIONS AND DYNAMIC RANGE REQUIREMENTS

APPLICATION	SIGNAL BANDWIDTH	DYNAMIC RANGE	ADC # BITS
Seismology	10Hz	146dB	24
Digital Audio	20kHz	100dB	18
Echo Cancelling	4kHz	84dB	14
Speech Processing	4kHz	74dB	12
V.32 Modems	4kHz	74db	14
Ultrasound	15MHz	60dB	10
Radar	5MHz	74dB	12
Broadband Receivers	5MHz	86dB	14

Figure 3.22

Figure 3.22 shows a number of applications which are suitable for DSP processing. The approximate bandwidth and dynamic range of the corresponding signal is given. There are actually two aspects to dynamic range: *dc* and *ac*. The dynamic range corresponds to the values given in Figure 3.22 (neglecting ADC static errors). AC dynamic range, on the other hand, is related to the harmonic distortion performance of the ADC. For instance, in a digital spectral analysis application, the harmonics of a fullscale sinewave input signal limits the system's

ability to resolve small signals in the presence of large signals. AC linearity is usually expressed in terms of harmonic distortion, or total harmonic distortion (THD). In a practical ADC, the number of bits may not be a good indicator of the harmonic distortion performance of the converter. AC dynamic range is less than that predicted by the SNR equation, $6.02N + 1.76\text{dB}$. For these reasons, the data sheet must include both *dc* and *ac* performance specifications in order for the user to make an intelligent selection for the application.

ADC STATIC TRANSFER CHARACTERISTICS

ADC STATIC PERFORMANCE SPECIFICATIONS

- Differential Non-Linearity (DNL)
- Integral Non-Linearity (INL)
- Missing Codes
- Gain Error
- Offset Error

Figure 3.23

The basic specifications which describe the static performance of an ADC are given in Figure 3.23. In the ideal transfer function for a 3 bit ADC (Figure 3.24), the analog input signal is on the horizontal axis and the digital output is on the vertical axis. The digital output of the ADC is valid over a range of input signal. The quantum of input for a given output code is called the *width* of the code. The ideal width is exactly 1 LSB (least significant bit), but, in practice, each code-width is different from its neighbors. The deviation in the code-widths from the ideal 1 LSB value is called differential non-linearity, or DNL. A 3-bit ADC with various errors is shown in Figure 3.25. Note that the code 100 is missing because of the large DNL

associated with the adjacent codes. Missing codes can produce oscillation and hunting in a closed-loop system, thus making this an important parameter to consider for ADC selection in this application.

Integral non-linearity, or INL, is usually measured with respect to the code centers. A straight line is drawn through the end-points, and the worst deviation of any code center from this ideal straight line is the INL as shown in Figure 3.25. In some cases, integral non-linearity is defined with respect to a *best-fit* straight line which is typically calculated using the least-squares method.

Gain and *offset* errors apply to all codes equally and are usually trimmed out in a system using fairly traditional techniques.

TRANSFER FUNCTION FOR IDEAL 3-BIT ADC




Figure 3.24

TRANSFER FUNCTION FOR NON-IDEAL 3-BIT ADC




Figure 3.25

DAC STATIC TRANSFER CHARACTERISTICS

The basic specifications which describe the static transfer characteristics of a DAC are

given in Figure 3.26.

DAC STATIC PERFORMANCE SPECIFICATIONS

- Differential Non-Linearity (DNL)
- Integral Non-Linearity (INL)
- Non-Monotonicity
- Gain Error
- Offset Error

Figure 3.26

The static transfer function for an ideal 3-bit DAC is shown in Figure 3.27. The digital input values are plotted on the horizontal axis and the corresponding analog output values on the vertical. Unlike an ADC, a DAC cannot have a missing code. There will be a discrete analog output voltage produced for each digital input code. Differential non-linearity is defined as the variation in the spacing between adjacent analog output values from the ideal 1 LSB value. Excessive DNL errors can result in non-monotonic conditions as shown in Figure 3.28. A DAC is said to be non-monotonic if an increase in the digital code input causes a decrease in the analog output value. Conversely, a DAC

is said to be monotonic if the slope of its transfer characteristic has the same sign over its entire range. Non-monotonic conditions can produce oscillations in a closed loop system; therefore, this specification is important in the selection of a DAC for such applications.

Integral non-linearity is defined as the worst case variation in any of the analog output values with respect to an ideal straight line drawn through the end points. As with an ADC, INL may also be defined with respect to a best-fit straight line.

Gain and offset definitions are similar to those for ADCs and affect each analog output value equally.

TRANSFER FUNCTION FOR IDEAL 3-BIT DAC




Figure 3.27

TRANSFER FUNCTION FOR NON-IDEAL 3-BIT DAC




Figure 3.28

ADC DYNAMIC PERFORMANCE

In order to be useful in most DSP applications, the ADC must have acceptable dc and ac performance characteristics. A listing of

the most important dynamic ADC characteristics is given in Figure 3.29.

ADC DYNAMIC SPECIFICATIONS

- Signal-to-Noise Plus Distortion (S/N + D) Ratio and Effective Number of Bits
- Peak Spurious, Peak Harmonic Content, and Spurious Free Dynamic Range (SFDR)
- Total Harmonic Distortion (THD)
- Full-Power Bandwidth (FPBW)
- Full-Linear Bandwidth
- Intermodulation Distortion (IMD)
- Aperture Delay Time and Aperture Jitter
- Transient Response
- Overvoltage Recovery

Figure 3.29

As we will see in a later section, there are a number of architectures which are suitable for DSP ADC designs, and most require a sample-and-hold amplifier (SHA) ahead of the actual converter as shown in Figure 3.30. Notable exceptions are flash converters and, in particular, sigma-delta converters. To fully characterize the dynamic performance of a SHA-ADC pair, they must be integrated onto the same chip, or at least offered as a complete functional unit. Otherwise, it is

almost impossible to determine the overall dynamic performance of the SHA-ADC combination from the specifications on the individual devices. The requirement for complete dc and ac characterization of ADCs has led to the introduction of *sampling* ADCs which have on-board SHAs. These converters eliminate the problems of interfacing SHAs to ADCs and provide users with complete dc and ac specifications.

SIGNAL-TO-NOISE RATIO AND EFFECTIVE BITS

As has been previously discussed, the signal-to-noise ratio specification is probably the most all-inclusive ac specification used in the industry today. Since it is common practice to include the effects of harmonic distortion in this measurement, S/N+D is defined as the ratio of the rms value of the measured input signal to the rms sum of all other spectral components below the Nyquist

frequency, including harmonics, but excluding dc. A typical plot of S/N+D is shown in Figure 3.31 for three high speed flash ADCs. The harmonic distortion performance of the AD9617 current feedback op amp is shown on the same graph for comparison. The SNR measurement can also be expressed in effective bits, or ENOBs, as is also shown in Figure 3.31.

ADC WITH TRACK-AND-HOLD




Figure 3.30

FLASH ADC AND OP AMP DYNAMIC PERFORMANCE




Figure 3.31

PEAK SPURIOUS, PEAK HARMONIC CONTENT, AND SPURIOUS FREE DYNAMIC RANGE (SFDR)

The peak spurious or peak harmonic component is the largest spectral component excluding the input signal and dc. This value is expressed in dB relative to the rms value of a fullscale input signal. The peak spurious

specification is also occasionally referred to a spurious free dynamic range (SFDR). A typical plot showing the peak spurious performance for the AD678 is shown in Figure 3.32.

PEAK SPURIOUS RESPONSE FOR AD678
AT 200 KSPS, NONAVERAGED 2048 POINT FFT




Figure 3.32

TOTAL HARMONIC DISTORTION (THD)

Total harmonic distortion (THD) is the ratio of the rms sum of the first six harmonic components to the rms value of a fullscale input signal and is expressed in a percentage or in dB. For input signals or harmonics

that are above the Nyquist frequency, the aliased component is used. Typical THD performance for the AD678 is shown in Figure 3.33.

TOTAL HARMONIC DISTORTION, FULL-POWER
BANDWIDTH, AND FULL LINEAR BANDWIDTH
FOR AD678




Figure 3.33

FULL-POWER BANDWIDTH

The full-power bandwidth (FPBW) of an ADC is that input frequency at which the amplitude of the reconstructed (using FFTs) *fundamental* is reduced by 3dB for a fullscale input. As can be seen from Figure 3.33, the full-power bandwidth of the AD678 is ap-

proximately 1MHz. In order to be meaningful, however, FPBW must be examined in conjunction with SNR, ENOB, and harmonic distortion in order to determine the true dynamic performance of the ADC at the FPBW frequency.

FULL-LINEAR BANDWIDTH

The full-linear bandwidth of an ADC is the input frequency at which the slewrate of the input sample-and-hold (SHA) is reached. At this point, the amplitude of the reconstructed sinewave has degraded by less than -0.1dB. Beyond this frequency, distortion of the sampled input signal increases significantly .

The AD678 ADC has been designed to optimize input bandwidth, allowing it to under-sample input signals significantly above the converter's Nyquist frequency. The full-linear bandwidth specification is 500kHz for the AD678 and is also shown in Figure 3.33.




Figure 3.34

INTERMODULATION DISTORTION (IMD)

Intermodulation distortion (IMD) occurs when the inputs consist of sinewaves at two frequencies, $F1$ and $F2$. Any device with nonlinearities will create distortion products, of the order $(m+n)$, at sum and difference frequencies of $mF1 \pm nF2$, where $m, n = 0, 1, 2, 3 \dots$. Intermodulation terms are those for which m or n is not equal to zero. For example, the second order terms are $(F1 + F2)$ and $(F1 - F2)$, and the third order terms are $(2F1 + F2)$, $(2F1 - F2)$, $(F1 + 2F2)$, and $(F1 -$

$2F2)$ (see Figure 3.34). The IMD products are expressed as the dB ratio of the rms sum of the measured input signals to the rms sum of the distortion terms. The two signals applied to the ADC are of equal amplitude and the peak value of their sum is -0.5dB from fullscale. The IMD products are normalized to a 0dB input signal. A typical IMD FFT plot for the AD678 is shown in Figure 3.35.




Figure 3.35

AC LINEARITY PLOTS USING HISTOGRAMS

For this measurement, a fullscale sine-wave is applied to the ADC, and a large number of samples are taken. The number of occurrences of each code is recorded on a histogram plot as shown in the top left-hand curve in Figure 3.36. In the case of a 12-bit converter, several million samples are required in order to achieve statistically significant results. The histogram should follow the ideal probability density distribution of a sine-wave, which is shown in the top right-hand curve in Figure 3.36. The histogram data is then normalized using the sine-wave

probability density function to obtain the DNL plot shown in the bottom curve of the figure. Integral non-linearity can be determined by compiling a cumulative histogram. The cumulative bin widths are the transition levels. Figure 3.37 shows an ac linearity plot obtained using the histogram method for the AD7870 12-bit 100kSPS ADC digitizing a 25kHz input signal at a 100kSPS rate. The absence of large spikes in the plot shows good differential linearity. More details of the mathematics involved is given in the AD7870 data sheet .

AC LINEARITY USING HISTOGRAMS

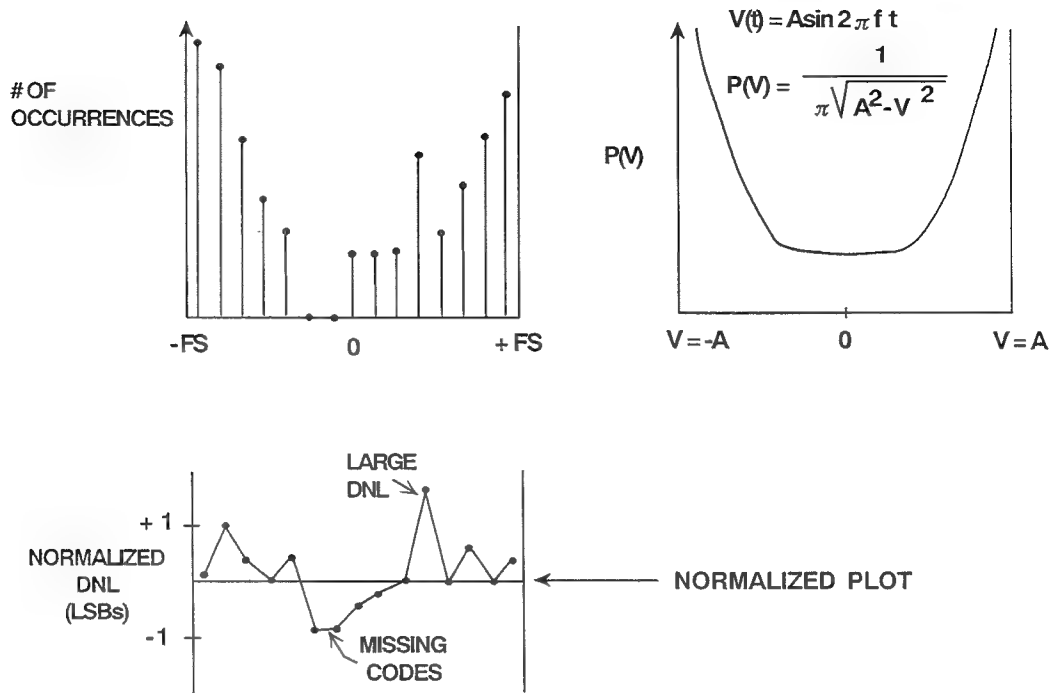


Figure 3.36

AC LINEARITY OF AD7870 12-BIT, 100kSPS ADC WITH 25 KHZ INPUT




Figure 3.37

APERTURE DELAY TIME (OR EFFECTIVE APERTURE DELAY TIME)

Aperture delay time (sometimes called aperture time) is the amount of time from the leading edge of the sampling clock until the ADC actually takes the sample (see Figure 3.38). This specification is important because it helps the user to know when to apply the sampling clock with respect to the

input signal timing. The variation or tolerance placed on this parameter from part to part is important in simultaneous sampling applications or other applications where the ADCs are required to track each other when processing dynamic signals.

MEASUREMENT OF EFFECTIVE APERTURE DELAY TIME




Figure 3.38

APERTURE JITTER

Aperture jitter is the sample-to-sample variation in the effective point in time at which the actual sample is taken as shown in Figure 3.39. These errors generally emanate from several sources. In a practical ADC, the sampling clock is often phase-modulated by some unwanted source; the source can be wideband random noise, power line noise, or digital noise due to poor layout, bypassing, and grounding techniques. The resulting error can be expressed in terms of an rms

time jitter. The corresponding rms voltage error caused by rms aperture jitter decreases the overall ADC signal-to-noise ratio. Phase jitter on the input sinewave can produce the same effect as jitter on the sampling clock. The SNR due exclusively to aperture jitter is plotted in Figure 3.40 as a function of full-scale sinewave input frequency for various values of aperture jitter. The equation for SNR due to aperture jitter is derived in Reference 1.

EFFECTS OF APERTURE JITTER




Figure 3.39

SIGNAL TO NOISE RATIO DUE TO APERTURE JITTER




Figure 3.40

TRANSIENT RESPONSE OR SETTLING TIME

The transient response (or settling time) of an ADC is the time required for the ADC to settle to rated accuracy after the application of a fullscale step input (see Figure 3.41). This specification is critical in applications where the ADC is being driven by an analog multiplexer as shown in Figure 3.42.

The multiplexer output can deliver a fullscale sample-to-sample change to the ADC input. If both the multiplexer and the ADC have not both sufficiently settled to the required accuracy, dc channel-to-channel crosstalk will result.

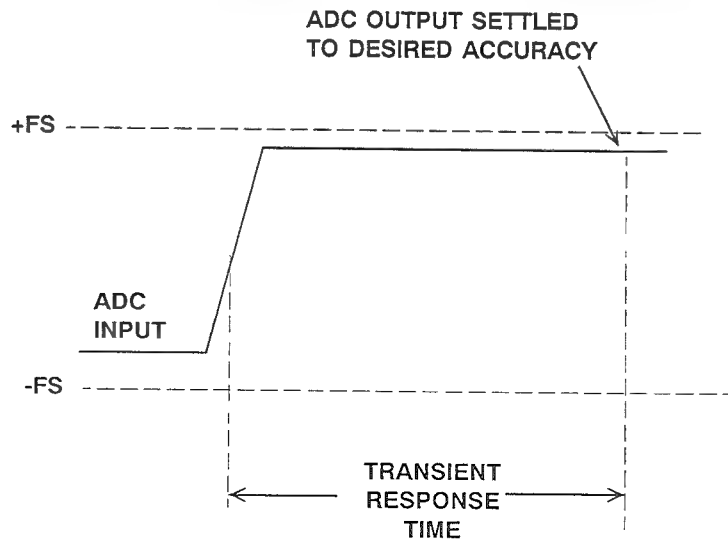
ADC TRANSIENT RESPONSE

Figure 3.41


TRADITIONAL DATA ACQUISITION SYSTEM USING ANALOG MULTIPLEXER

Figure 3.42

OVERVOLTAGE RECOVERY

Overvoltage recovery time is defined as that amount of time required for the ADC to achieve a specified accuracy, measured from the time the overvoltage signal re-enters the converter's range, as shown in Figure 3.43. This specification is usually given for a signal which is 50% outside the ADC's input range.

Needless to say, the ADC should act as an ideal limiter for out-of-range signals and should produce either the positive fullscale code or the negative fullscale code during the overvoltage condition. Some converters provide over- and under-range flags to allow gain-adjustment circuits to be activated.

ADC OVERVOLTAGE RECOVERY

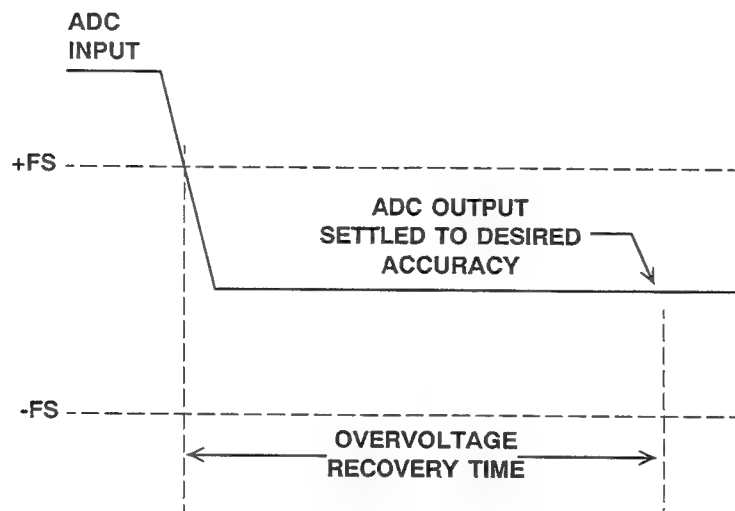


Figure 3.43

DAC DYNAMIC PERFORMANCE

Since most DSP applications involve the eventual reconstruction of a dynamic analog signal, ac performance of DACs has become

as important as ADC performance. Key DAC ac performance characteristics are given in Figure 3.45.

DAC DYNAMIC SPECIFICATIONS

- Settling Time
- Glitch Impulse Area
- Harmonic Distortion
- Signal-to-Noise Ratio
- Audio-Specific Specifications

Figure 3.44

DAC SETTLING TIME WAVEFORM

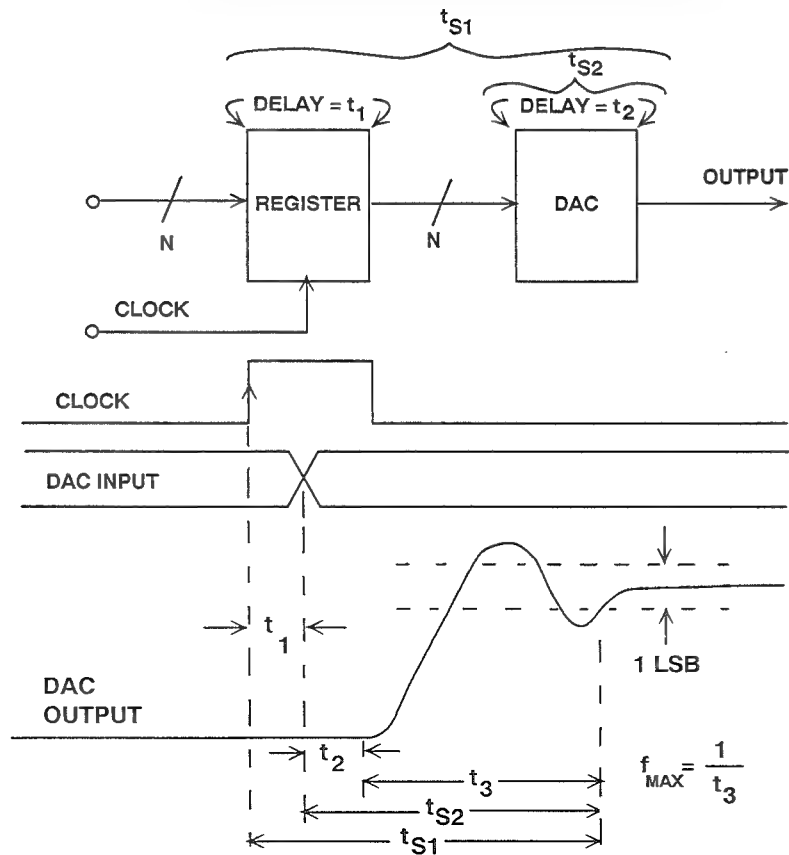


Figure 3.45

SETTLING TIME

Settling time of a DAC is traditionally defined as the time from the digital input transition (usually measured from the 50% point) until the DAC output settles to within a certain error band (usually $1/2$ LSB) which is centered around the final value. As shown in Figure 3.45, a portion of the settling time may be due to a fixed propagation delay through the switches. If the DAC has a set of input latches or registers, the settling time should be measured from the 50% point of the latch strobe or register clock. Fullscale DAC settling time is measured for a digital input transition from 000...0 to 111...1. Midscale settling time is measured for a

digital transition from 011...1 to 100...0 or 100...0 to 011...1.

It is entirely correct to define DAC settling time with respect to the output alone as shown in Figure 3.46. Settling time is measured from the time the output leaves a $\pm 1/2$ LSB error band centered around the initial value until the time the output remains within a $\pm 1/2$ LSB error band centered around the final value. The maximum DAC update rate allowable for $\pm 1/2$ LSB fullscale settling time then becomes $f_{\max} = 1/t_s$. Faster update rates can be used if sample-to-sample changes in the DAC input are limited to values less than fullscale.

SETTLING TIME DEFINED WITH RESPECT TO DAC OUTPUT




Figure 3.46

GLITCH IMPULSE AREA


Glitch impulse area is best understood by examining the waveform shown in Figure 3.47. DAC glitches occur because of digital input logic skew and unequal propagation delays through the DAC switches (a noteworthy exception to this is the sigma-delta DAC architecture to be discussed later in this seminar). The glitches are usually the largest at the midscale transition because all bits in the DAC are changing at this point. The glitch produced by the 011...1 to 100...0 transition is usually different from that produced by the 100...0 to 011...1 transition, so each must be analyzed. Glitch impulse area is simply the area of a particular glitch, and is usually measured in the units of pV-sec, therefore the fullscale output voltage of the DAC must be known in order to make meaningful comparisons between DACs. The term *glitch energy* is incorrect since the unit pV-sec is *not* a measure of energy.

From Figure 3.47 it is clear that there are six possible glitch impulse areas to deal with.

There are two glitch impulses associated with each transition. Their respective areas are designated 1,2,3, and 4. In addition, it is also useful to consider the *net glitch impulse* area associated with each of the two transitions. There are, respectively, AREA 1 - AREA 2, and AREA 3 - AREA 4. When examining the glitch impulse area specification on a DAC data sheet, it is therefore clear that there is much room for confusion unless a considerable amount of clarification is provided by the manufacturer.

Glitch impulse area remains constant regardless of filtering. Fast settling time specifications do not always imply low glitch impulse areas. The desirable situation is for the DAC to have a net glitch impulse area of zero for each of the two transitions, i.e., AREA 1 - AREA 2 = AREA 3 - AREA 4 = 0. In the ideal case, of course, each of the four areas would be zero.

GLITCH IMPULSE WAVEFORMS



GLITCH IMPULSE AREAS: AREA 1
 AREA 2
 AREA 3
 AREA 4

NET GLITCH IMPULSE AREAS: | AREA 1 - AREA 2 |
 | AREA 3 - AREA 4 |

Figure 3.47

HARMONIC DISTORTION

Because the net glitch impulse area is code-dependent, it will produce harmonics when the DAC is reconstructing a sinewave. A net midscale glitch occurs twice during a single cycle of the reconstructed sinewave (at each zero crossing) and, therefore, will produce a second harmonic of the sinewave as shown in Figure 3.48. Note that higher order harmonics of the sinewave which alias back into the Nyquist bandwidth are not filterable. It is difficult to predict the harmonic distortion caused by a specified net glitch impulse area, therefore, both specifications are required to adequately evaluate the dynamic performance of a reconstruction DAC.

Total harmonic distortion (THD) can be measured using DSP techniques as shown in Figure 3.49 for the AD1860 18-bit audio

DAC. The DAC is driven with an 18-bit digital sinewave having a frequency of 990.5Hz, and the DAC update rate is 176.4kHz. The DSP digitizes 4096 samples of the output test waveform, incorporating 23 complete cycles of the sinewave. A 4096 point FFT is performed on the results of the test. The total harmonic distortion and the SNR is then calculated from the FFT results. The notch filter prevents the large-amplitude fundamental component at 990.5Hz from entering the digitizer, thereby allowing the entire digitizer range to be dedicated to processing the noise and harmonic components. Figure 3.50 shows a typical THD + noise plot for both a fullscale input and a -20dB input. It should be noted that neither a deglitcher nor an MSB trim are used in these measurements.

EFFECTS OF DAC GLITCHES




Figure 3.48

FFT TESTING OF AD1860 18-BIT AUDIO DAC

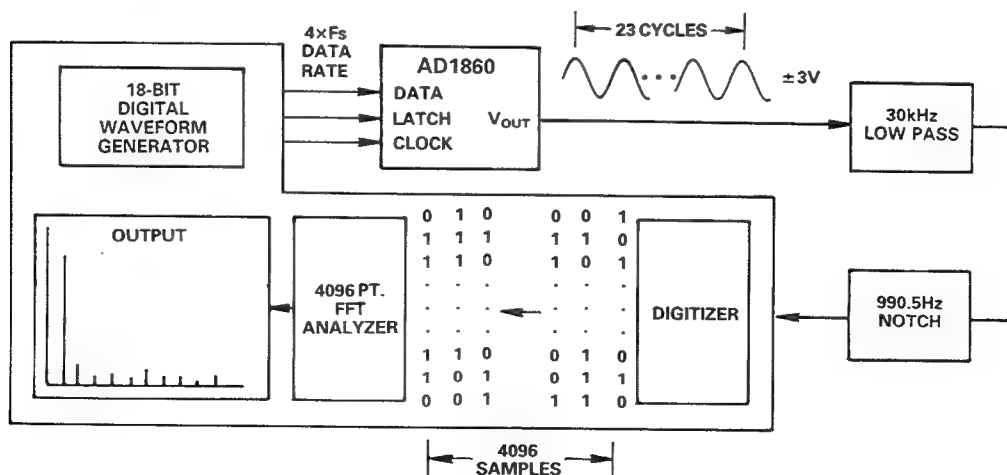


Figure 3.49

FFT-MEASURED TOTAL HARMONIC DISTORTION FOR AD1860 18-BIT AUDIO DAC




Figure 3.50

DEGLITCHING DACs USING SHAs

SHAs can be used to deglitch DACs as shown in Figure 3.51. Just prior to latching new data into the DAC, the SHA is put into the hold mode so that the DAC switching glitches are isolated from the output. The

switching transients produced by the SHA are code-independent and occur at the update frequency, hence, they are easily filterable.

SIN(X)/X FREQUENCY ROLLOFF EFFECT

The output of a reconstruction DAC can be visualized as a series of rectangular pulses whose width is equal to the reciprocal of the update rate as shown in Figure 3.52. Note that the reconstructed signal is down 3.92dB

at the Nyquist limit with respect to the low frequency value. An inverse $\sin(x)/x$ filter is sometimes placed after the DAC to correct for this effect.

SHA USED AS DEGLITCHER

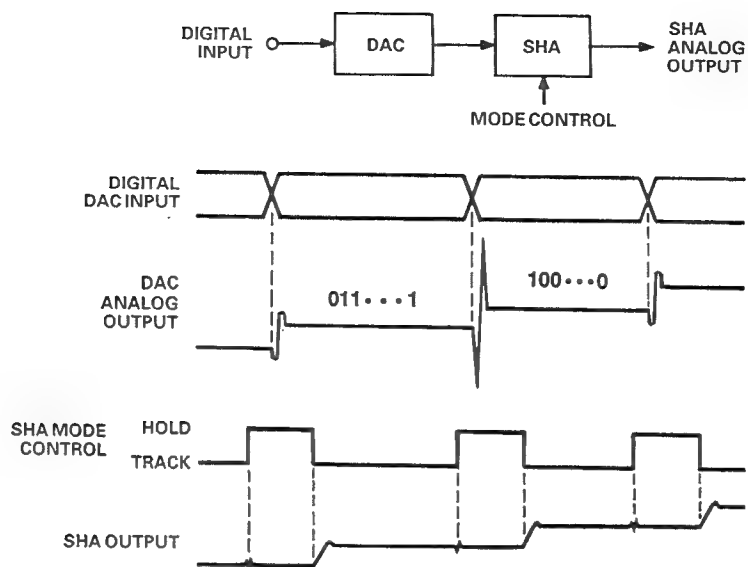


Figure 3.51

SIN X/X ROLLOFF




Figure 3.52

SWITCHED CAPACITOR FILTERS

Signals were once filtered entirely in the continuous analog domain by configurations of passive components (typically inductors, resistors, and capacitors). Later, active filters, with op amps for buffering and gain, provided filter designers with additional flexibility and performance, but still operated continuously on analog signals. DSP led to stable and flexible discrete-time digital filters, where sampled analog signals are processed entirely by numerical calculations with filtering algorithms—some of which cannot be realized with continuous-time analog filters.

Switched-capacitor filters (SCFs) are an intermediate class, combining both continuous- and discrete-time aspects. They are usually implemented using CMOS switches and capacitors to simulate the behavior of resistors, therefore, many filter architectures can be realized entirely by a monolithic device without the need for external components. SCFs are particularly useful for voice and audio bandwidth signal applications in conjunction with DSP technology. Since SCFs are *sampling* devices, all the concepts of discrete time sampling apply to their use: Nyquist's theorem, aliasing, etc.

FILTERING TECHNIQUES

- Crystals, SAWs
- Passive Components (R, L, and C)
- Active Filters (R, C, and Op-Amps)
- Switched Capacitor Filters (CMOS Switches and Capacitors Replace Resistors)
- Digital Filters (Numerical Realizations Which May Have No Analog Counterpart)

Figure 3.53

The fundamental concept of a switched capacitor acting as a resistor can best be understood from the concept of charge transfer as shown in Figure 3.54. If the capacitor is switched from V_1 to V_2 , an instantaneous charge transfer occurs, $\Delta Q = C(V_1 - V_2)$, either into or out of V_2 . This assumes that C has no series resistance, and that V_1 and V_2 are ideal voltage sources. If the switch is thrown back and forth at a clock frequency, f_s (having a period T), then an average current, i , flows between V_1 and V_2 having a value $i = \Delta Q/T = C\Delta V/T$. The equivalent resistance, "R", that would give the same average current is given by :

$$\text{"R"} = \Delta V/i = T/C = 1/(Cf_s).$$

In an integrated circuit, the single-pole double-throw switch is implemented using CMOS switches driven by a non-overlapping two-phase clock as shown in Figure 3.55. A requirement for this technique to work is that the switches have very low on resistance and very high off resistance. This is precisely what CMOS technology offers.

Using this SC resistor-equivalent, many conventional passive and active filter configurations can be realized. Figure 3.56 shows a single-pole passive RC filter and its SCF equivalent. The -3dB frequency of the RC filter is $1/(2\pi R_1 C_1)$. For the SCF version,

$$f_{3dB} = f_s C_1/(2\pi C_2).$$

Note that for the SCF version, the *bandwidth depends on the sampling rate and the ratio of the capacitor values*. A major assumption which must be made is that $f_s \gg f_{3dB}$ (typically 50 to 100) to minimize the effects of time-sampling and charge-sharing. Using the SCF concept, critical frequencies are therefore determined by capacitor ratios and the sampling clock frequency, both of which can be made precise and drift free.

Audio and voiceband filtering with SC filters can greatly reduce passive component

physical size. To implement audio filters, a resistance on the order of $10M\Omega$ is required if a monolithic capacitor of reasonable size ($\sim 10pF$) is to be used. This value of resistance is easily achieved by switching a $1pF$ capacitor at a $100kHz$ rate, requiring a silicon area of approximately $0.01mm^2$. If the $10M\Omega$ resistor were implemented using polysilicon or diffusion, the area required would be at least 100 times larger.

SWITCHED CAPACITOR "RESISTOR"




Figure 3.54

CMOS IMPLEMENTATION OF SWITCHED CAPACITOR




Figure 3.55

SC EQUIVALENT OF PASSIVE RC NETWORK




Figure 3.56

SWITCHED CAPACITOR FILTER ADVANTAGES

- Filter Bandwidths Proportional to Capacitance Ratios
Not Absolute Values
- Filter Bandwidths Variable with Clock Frequency
- Defined Like Classic Analog Filters
- Low Values of Capacitance Required for Audio Frequencies:
1pF Capacitor Switched at 100kSPS = $10\text{M}\Omega$ “Resistance”
- SCFs Ideally Suited to DSP CMOS Processes

Figure 3.57

By using SC resistors in conjunction with other capacitors and op amps, it is possible to realize many of the circuit configurations used in conventional RC active filters. Unlike digital filters, SC filters may be defined exactly like analog filters. A first-order continuous-time active lowpass RC filter and its SC counterpart are shown in Figure 3.58.

Since they sample analog signals, SC filters must usually be preceded by a continuous-time antialiasing prefilter to eliminate spectral components above the Nyquist frequency. Since the SC filter sampling rate is usually much higher than its passband, a single or double pole RC filter is usually sufficient for this purpose.

Differential amplifiers are often used in analog circuits to achieve good common mode rejection of unwanted signals such as power line noise, etc. The same principles can be used in designing switched capacitor filters. Figure 3.59 shows an active differential integrator and its switched capacitor equivalent.

In addition to providing good CMRR to noise, the differential configuration also provides common mode rejection to the transients caused by the operation of the switches. Switched capacitor integrators are often used in the modulator circuits of a Sigma-Delta ADCs as will be discussed later in this seminar.

Switched capacitor filters are subject to several limitations and error sources. Their usefulness is limited to frequencies in the audio bandwidth, since sampling rates greater than a few hundred kilohertz cannot be readily achieved with current CMOS technology. The switched capacitors and op amps introduce random noise, and leakage currents can produce offset errors. Clock feedthrough from the switches themselves can produce synchronous errors. Finally, since SCFs are sampling devices, large oversampling ratios are usually required in order to prevent errors due to aliasing.

FIRST ORDER ACTIVE LOWPASS RC FILTER AND SCF EQUIVALENT


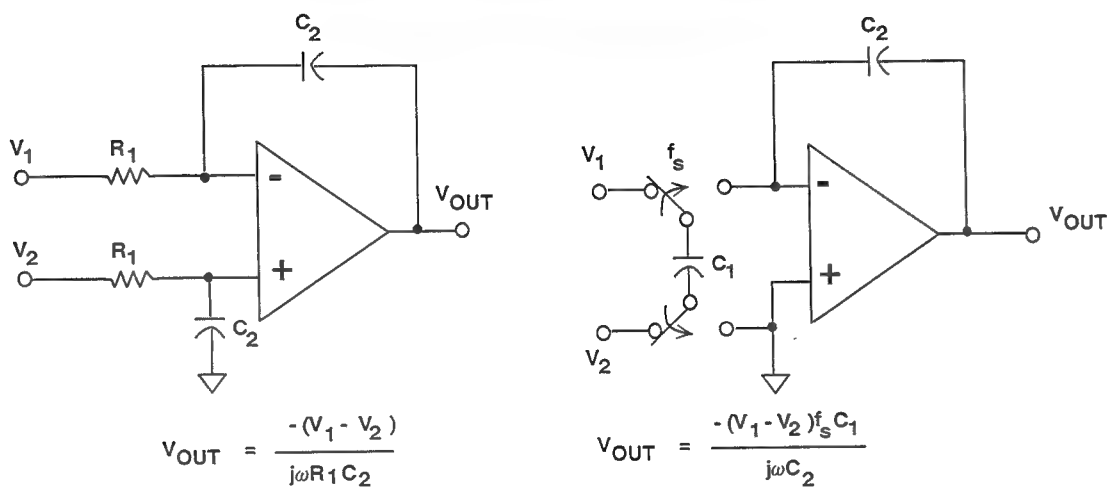


Figure 3.58

ACTIVE DIFFERENTIAL INTEGRATOR AND SCF EQUIVALENT



$$V_{OUT} = \frac{-(V_1 - V_2)}{j\omega R_1 C_2}$$

$$V_{OUT} = \frac{-(V_1 - V_2) f_s C_1}{j\omega C_2}$$

Figure 3.59

SWITCHED CAPACITOR LIMITATIONS AND ERROR SOURCES

- Limited to Lower Frequencies
- Noise, Offset, and Distortion
- Clock Feedthrough from Switches
- Must Obey the Laws of Nyquist (Requires Antialiasing Filter)

Figure 3.60

REFERENCES

1. Frederic J. Harris, "On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform", **IEEE Proceedings**, Vol. 66, No. 1, Jan. 1978, pp. 51-83.
2. Joey Doernberg, Hae-Seung Lee, David A. Hodges, "Full Speed Testing of A/D Converters", **IEEE Journal of Solid State Circuits**, Vol. SC-19, No. 6, Dec. 1984, pp. 820-827.
3. James R. Andrews, Barry A. Bell, Norris S. Nahman, and Eugene E. Baldwin, "Reference Waveform Flat Pulse Generator", **IEEE Transactions on Instrumentation and Measurement**, Vol. IM-32, No. 1, March 1983, pp. 27-32.
4. Brendan Coleman, Pat Meehan, John Reidy and Pat Weeks, "Coherent Sampling Helps When Specifying DSP A/D Converters", **EDN**, October 15, 1987, pp. 145-152.
5. Howard K. Schoenwetter, "A Programmable Voltage Step Generator for Testing Waveform Recorders", **IEEE Transactions on Instrumentation and Measurement**, Vol. IM-33, No. 3, Sept. 1984, pp. 196-200.
6. Robert W. Ramirez, **The FFT: Fundamental and Concepts**, Prentice-Hall, 1985.
7. R.B. Blackman and J.W. Tukey, **The Measurement of Power Spectra**, Dover Publications, New York, 1958.
8. James J. Colotti, "Digital Dynamic Analysis of A/D Conversion Systems Through Evaluation Software Based on FFT/DFT Analysis", **RF Expo East 1987 Proceedings**, Cardiff Publishing Co., pp. 245-272.
9. Sid Kaufman, "Multistage Error Correcting A/D Converters", **Electronic Products**, April 18, 1983, pp. 103-110.
10. **HP Journal**, Nov. 1982, Vol. 33, No. 11
11. HP Product Note 5180A-2.
12. **HP Journal**, April 1988, Vol. 39, No. 2.
13. **HP Journal**, June 1988, Vol. 39, No. 3.
14. Dan Sheingold, Editor, **Analog-to-Digital Conversion Handbook, Third Edition**, Prentice-Hall, 1986.
15. W.R. Bennett, "Spectra of Quantized Signals", **Bell System Technical Journal**, No. 27, July 1948, pp. 446-472.
16. G.A. Gray and G.W. Zeoli, "Quantization and Saturation Noise Due to Analog-Digital Conversion", **IEEE Transactions on Aerospace and Electronic Systems**, Jan. 1971, pp. 222-223.

17. M.J. Tant, **The White Noise Book**, Marconi Instruments, July 1974.
18. W.A. Kester, "PCM Signal Codecs for Video Applications", **SMPTE Journal**, No. 88, November 1979, pp. 770-778.
19. Lawrence Rabiner and Bernard Gold, **Theory and Application of Digital Signal Processing**, Prentice-Hall, 1975.
20. Matthew Mahoney, **DSP -Based Testing of Analog and Mixed-Signal Circuits**, , IEEE Computer Society Press, Washington, D.C., 1987.
21. IEEE Trial-Use Standard for Digitizing Waveform Recorders, No. 1057-1988.
22. Richard J. Higgins, **Digital Signal Processing in VSLI**, Prentice-Hall, 1990.
23. **High Speed Design Seminar**, Analog Devices, 1990.
24. M. S. Ghausi and K. R. Laker, **Modern Filter Design: Active RC and Switched Capacitors**, Prentice Hall, 1981.
25. Brodersen, Gray, and Hodges, *MOS Switched-Capacitor Filters*, **Proc. IEEE Vol. 67**, January 1979, pp. 61 - 65.
26. J. T. Caves, et. al., *Sampled Analog Filtering Using Switched Capacitors as Resistor Equivalents*, **IEEE J. Solid State Circuits Vol. SC-12**, pp. 592 - 599.

SECTION IV

ADCs FOR DSP APPLICATIONS

ADCs FOR DSP APPLICATIONS

- **SUCCESSIVE APPROXIMATION ADCs**
- **FLASH ADCs**
- **SUBRANGING ADCs**
- **INTEGRATING (DUAL SLOPE) ADCs**



SECTION IV

ADCs FOR DSP APPLICATIONS

The trend in ADCs for DSP applications is to integrate the sample-and-hold function with the ADC, thus producing what is commonly referred to as a *sampling* converter. This greatly simplifies system designs by eliminating the need for interfacing the SHA to the ADC, a process which involves tricky timing issues as well as a certain amount of

optimization for best performance. A sampling ADC can be completely specified not only in terms of traditional dc parameters, but also in terms of ac performance parameters such as SNR, effective bits, THD, etc. In the following section, we examine several popular architectures used to implement the ADC function.

ADCs for DSP APPLICATIONS

- Most are *Sampling* ADCs Containing on-chip SHA, or Sigma-Delta Converters
- No Need for Customer to Interface SHA to ADC Encoder
- Complete DC and AC Specifications Provided

Figure 4.1

SUCCESSIVE APPROXIMATION ADCs

The successive approximation architecture shown in Figure 4.2 has been extremely popular in the industry primarily because it combines relatively high resolution and speed with low cost. The building blocks for the encoder portion of the successive approximation ADC consist of a comparator, DAC, and control logic (successive approximation register, or SAR). The overall static accuracy is primarily determined by the DAC which can be laser trimmed at the wafer level.

The analog input drives one input of the comparator, while the DAC output is connected to the other input. The conversion technique consists of comparing the unknown input against a precise voltage or current generated by the DAC. The input of the DAC is the digital number at the ADC's output. The conversion process is strikingly similar to a weighing process using a chemist's balance, with a set of N binary weights (e.g., 1/2 lb, 1/4 lb, 1/8 lb, 1/16 lb, etc.) for unknowns up to 1 lb.

After the conversion command is applied, and the converter has been cleared, the DAC's MSB output (1/2 fullscale) is compared with the input. If the input is greater than the MSB, it remains ON (i.e., "1" in the output register), and the next bit (1/4 FS) is tried. If the input is less than the MSB, it is turned OFF (i.e., "0" in the output register), and the next bit is tried. If the second bit doesn't add enough weight to exceed the input, it is left ON ("1") and the third bit is tried. The process continues in order of descending bit weight until the last bit has been tried. When this process is completed, the conversion complete line changes state to indicate that the contents of the output register now constitute a valid conversion. The contents of the output register form a binary word corresponding to the input signal's magnitude. In most successive approximation ADCs, the output data is also available in serial format.

SUCCESSIVE APPROXIMATION A/D CONVERTER ENCODER

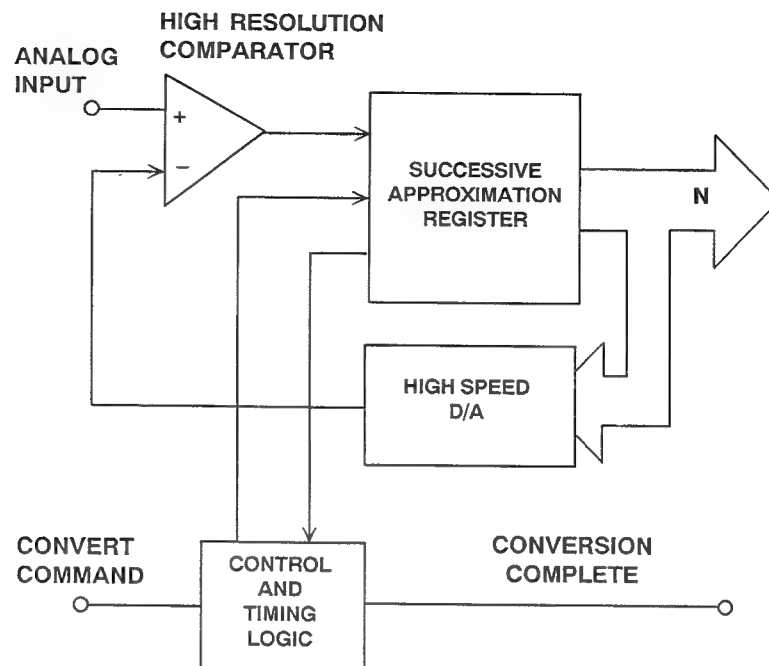


Figure 4.2

SAR ADC WITH SHA




Figure 4.3

Each of the bit decisions requires a clock period. An N-Bit converter will have N clock periods (plus an initialization period). Thus, the minimum conversion time of the ADC will be determined by the maximum allowable clock frequency and the number of bits. This frequency is limited by several factors: SAR clock-to-data-output delay, DAC settling time, and SAR input data setup time.

During the conversion time, it is important that the analog input signal be held constant to avoid errors. This usually requires that a SAR ADC encoder be preceded by an appropriate SHA if dynamic signals

are to be digitized. Sampling ADCs integrate the entire function shown in Figure 4.3 on a single chip. An example of this integration is the AD1674, a 12 bit 100kSPS ADC with on-chip SHA. This device is packaged with industry-standard AD574 pinouts and has complete dc and ac specifications. The AD7870/AD7875/AD7876 is a family of 12 bit 100kSPS low-power CMOS ADCs which also has the on-chip SHA. These devices offer complete ac and dc specifications as well as flexible input voltage ranges and easy parallel or serial interfacing to popular DSP processors.

FLASH ADCs

Recent advances in VLSI process technology and design techniques have made flash ADCs with up to 10-bits of resolution practical. As well as offering high sampling rates for digitizing video signals (usually without requiring a SHA), flash converters are often used as building blocks for higher resolution ADCs. A block diagram of a typical flash converter is shown in Figure 4.4. The analog

input signal to be digitized is applied simultaneously to $2^N - 1$ latched comparators, where N is the number of bits. The reference voltage input for each comparator is derived from a resistive voltage divider string. The reference voltage for each comparator is one LSB (least significant bit) higher than the comparator immediately below it.

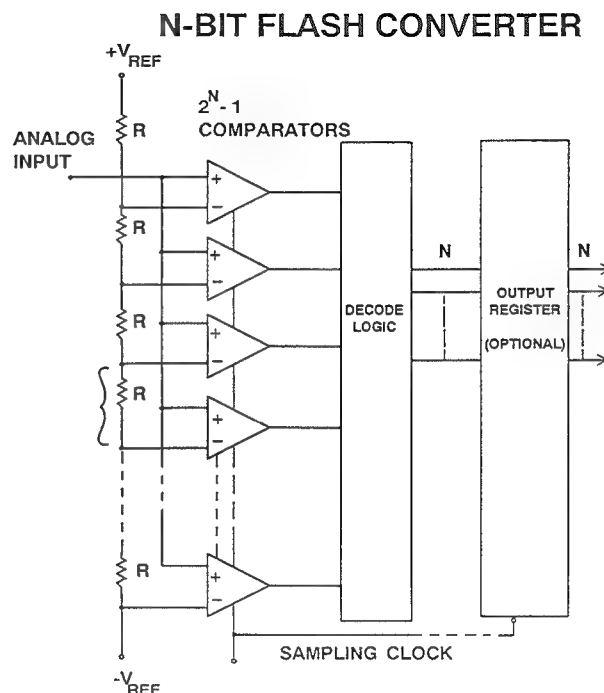


Figure 4.4


When an analog signal is present at the input of the comparator bank, all comparators which have a reference voltage below the level of the input signal will assume a logic "1" output. The comparators which have their reference voltage above the input signal will assume a logic "0" output. The result is often referred to as a *thermometer* code, and is applied to a stage of decoding logic. This decoding can be accomplished in a variety of ways (such as a simple priority encoder) and ultimately results in the formation of the digital output word. As shown in the diagram, the binary output of the decoding logic often drives an on-chip output latch.

Flash converters are inherently sampling ADCs and usually require no SHA, but they do require careful selection of the drive

amplifier because of their low input resistance, high bandwidth and rather large (and non-linear) input capacitance. Typical effective bit performance of several flash converters is shown in Figure 4.5 along with the harmonic distortion performance of an appropriate drive amplifier.

The AD9060 10 bit, 75MSPS flash converter meets the exacting requirements of HDTV systems as well as other instrumentation applications. The AD9038 8 bit, 300MSPS converter provides a demultiplexed output to simplify the interface with high speed memories. The AD9058 dual 8 bit 50MSPS ADC offers an effective solution where matching is important (In-phase and quadrature radar receivers, and ultrasound imaging) or PC board space is at premium.

FLASH ADC AND OP AMP DYNAMIC PERFORMANCE



FLASH ADC	RESOLUTION	SAMPLING RATE
AD9020	10 BITS	40 MSPS
AD9060	10 BITS	60 MSPS
AD9028/9038	8 BITS	250 MSPS
AD9006/9016	6 BITS	400 MSPS

Figure 4.5

SUBRANGING ADCs

A block diagram of an 8-bit subranging ADC based upon two 4-bit flash converters is shown in Figure 4.6. Although 8-bit flash converters are readily available at high sampling rates, this example will be used to illustrate the theory. The conversion process is done in two steps. The first four significant bits (MSBs) are digitized by the first flash (to better than 8-bits accuracy), and the 4-bit binary output is applied to a 4-bit DAC (better than 8-bit accurate). The DAC output is subtracted from the held analog input, and the resulting residue signal is amplified and applied to the second 4-bit flash converter by the summing amplifier. The outputs of the two 4-bit flash converters are then combined into a single 8-bit binary output word. If the residue signal doesn't exactly fill the range of the second flash converter, non-linearities and missing codes will result.

Modern subranging ADCs use a technique called *digital correction* to eliminate problems associated with the architecture of Figure 4.6. A block diagram of a 12-bit digitally corrected subranging (DCS) ADC is shown in Figure 4.7. Note that a 5-bit and an 8-bit flash converter have been used to achieve an overall 12-bit output. If there were no errors, the 5-bit "residue" signal applied to the 8-bit flash converter by the summing amplifier would never exceed one-half of the range of the 8-bit flash. The extra range in the second flash converter is used in conjunction with the error correction logic (usually just an adder) to correct the output data for most of the errors inherent in the traditional uncorrected subranging converter architecture previously discussed.

4

8-BIT SUBRANGING A/D CONVERTER

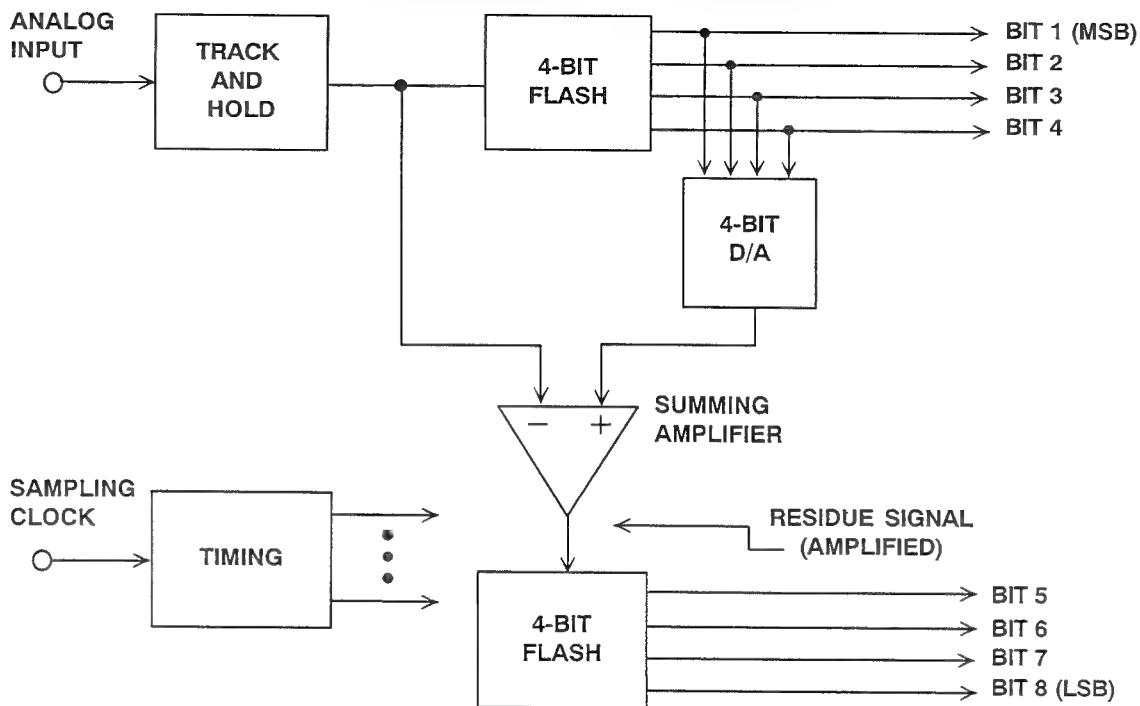


Figure 4.6

12-BIT SUBRANGING A/D CONVERTER - DCS

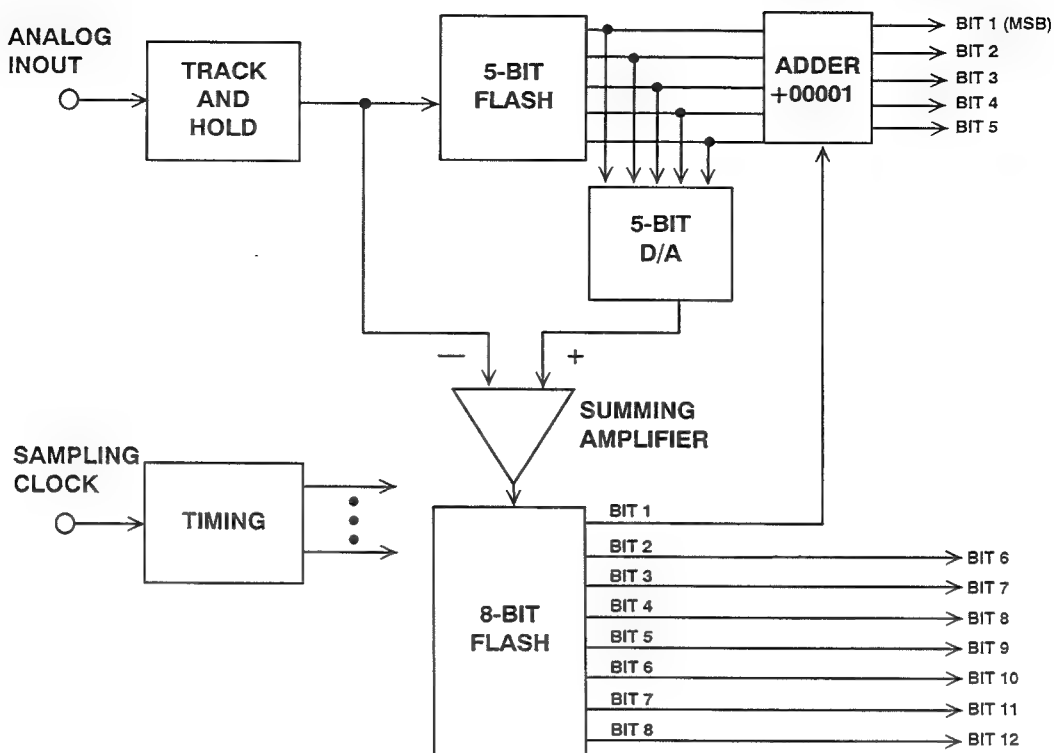


Figure 4.7

AD1671 12-BIT, 1 MSPS DIGITALLY CORRECTED SUBRANGING ADC

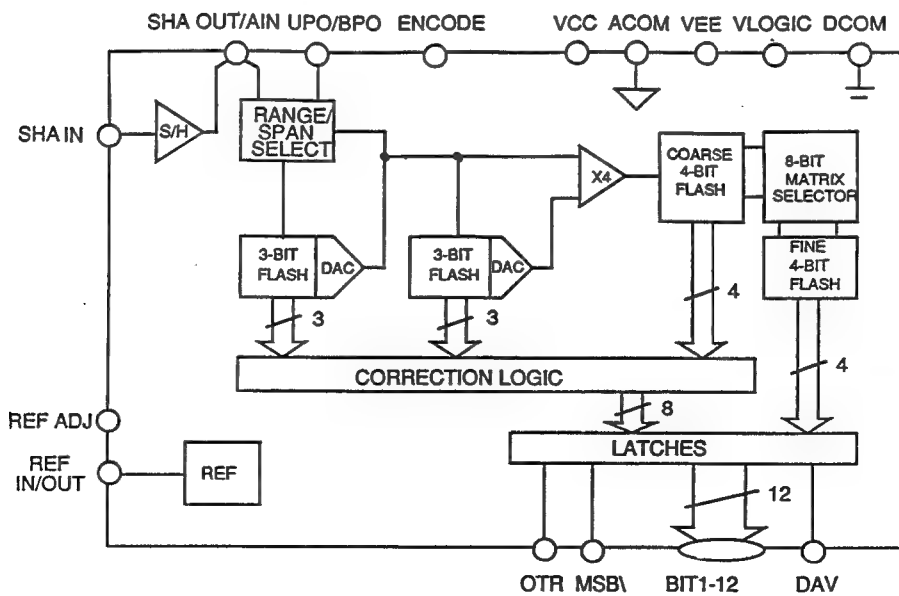


Figure 4.8

The AD1671 ADC uses a four-step DCS architecture as shown in Figure 4.8. Upon receipt of an ENCODE command, the SHA goes into the hold mode. After sufficient time for settling has elapsed, the first 3-bit flash converts the analog input voltage. The 3-bit residue is passed to a correction logic register and a segmented current output DAC. The DAC output is connected through a resistor (within the Range/Span Select Block) to AIN. A residue voltage is created by subtracting the DAC output from AIN, which is less than one-eighth of the fullscale analog input. The second flash has an input range that is configured with one bit of overlap with the previous DAC. The overlap allows for errors during the flash conversion. The first residue voltage is connected to the second 3-bit flash and to the noninverting input of a high speed, differential, gain-of-four amplifier. The second flash result is passed to the correction logic register and to the second segmented current output DAC. The output of the second DAC is connected to the inverting input of the differential amplifier. The differential amplifier output is connected to a two-step backend 8-bit flash. This 8-bit flash consists of coarse and fine flash converters. The result of the coarse 4-bit flash converter,

also configured to overlap one bit of DAC 2, is connected to the correction logic register and selects one of 16 resistors from which the fine 4-bit flash will establish its span voltage. The fine 4-bit flash is connected directly to the output latches.

The AD679 14-bit 100kSPS ADC shown in Figure 4.9 uses a *recursive subranging* architecture which makes multiple passes through a single 4-bit flash converter. The 4-bit flash first produces a 4-bit representation of the analog input. This value is reconstructed through the DAC, and the difference between this and the actual input is then amplified (to take full advantage of the dynamic range of the 4-bit flash), and the whole cycle then repeats itself. After 5 cycles, the result is presented at the digital output. A 1-bit overlap between cycles serves as error-correction.

The AD7886 is a 12 bit, 650kSPS low-power CMOS subranging ADC with on-chip SHA which is completely specified for both dc and ac performance. The AD9014 is a 12 bit, 10MSPS converter optimized for use in broadband receiver applications where spurious-free dynamic range is a key specification.

AD679 14-BIT, 100 kSPS
RECURSIVE SUBRANGING ADC

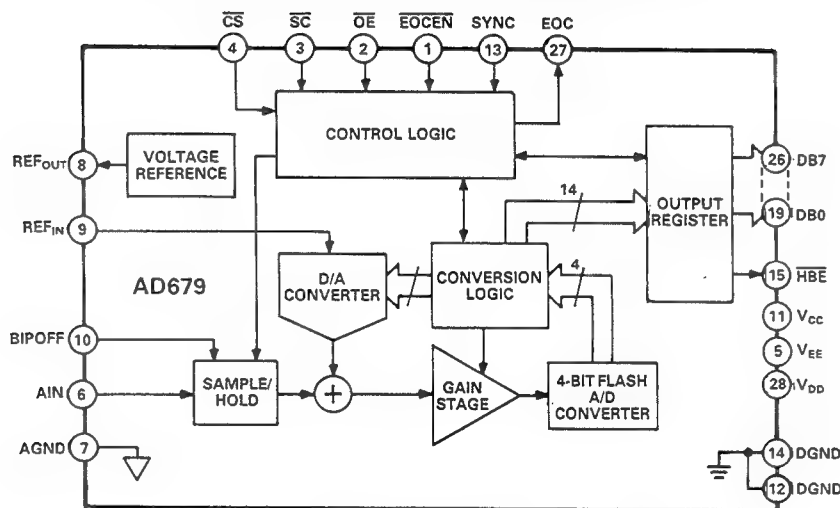


Figure 4.9

INTEGRATING (DUAL SLOPE) ADCs

Integrating ADCs are used for slow, precise measurements such as in digital voltmeters and many applications involving slow transducers, especially when they drive a display. An example of this type of converter is the high performance AD1175 22-bit 10, 20HZ ADC. The basic integrating ADC block diagram is shown in Figure 4.10 and consists of an integrator, a comparator, an up/down counter, a clock, and control logic.

Before the conversion starts, the integrator is held in the discharged state, i.e. R is connected to ground via S1 and C is shorted by S2. At the start of the conversion, R is connected to the unknown input voltage via S1, and S2 opens, allowing C to charge. The fixed integration time is controlled by the clock and the counter. At the end of the

integration period, S1 connects a known reference voltage to R, and the capacitor is discharged. The amount of time required for the capacitor to discharge is measured by the counter. Since the integrating capacitor, the resistor, and the clock frequency are unchanged during the charge and discharge cycles, the ratio of the charge and discharge times is equal to the ratio of the reference voltage to the unknown input voltage. The conversion accuracy is unaffected by the absolute values of R, C, or the clock frequency, and any noise on the input signal is integrated for the whole sampling period. A diagram showing the charging and discharging action of the integrating ADC is shown in Figure 4.11.

INTEGRATING (DUAL SLOPE) ADC

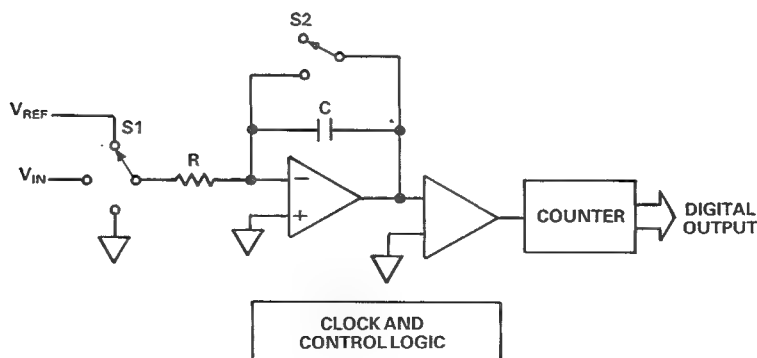


Figure 4.10

CHARGE/DISCHARGE WAVEFORM OF INTEGRATING ADC




Figure 4.11

Leakage and offset in the integrator and hysteresis and offset in the comparator affect the conversion accuracy, and most practical integrating ADCs perform a conversion with the integrator input grounded to compute the errors due to these effects. These errors are then subtracted from the result of an actual conversion. The nature of an integrating ADC is such that there is virtually no differential nonlinearity, and it does not suffer from missing codes.

The sampling period, T , of an integrating ADC is fixed. If the frequency of any ripple on the input to the ADC is an integral multiple of $1/T$, a whole number of cycles will occur during each integration period, and the

net contribution to the charge in the integrator due to ripple will be zero. Thus, an integrating ADC has near infinite rejection of input frequencies n/T , where n is an integer. By selecting T to be the period of the power line frequency (16.667ms for 60Hz and 20ms for 50Hz), line ripple on the input signal will be disregarded. The normal mode response for the integrating ADC with a conversion period T is shown in Figure 4.12.

As will be discussed in Section VI, sigma-delta ADC technology offers an attractive alternative to the integrating ADC, especially if the function is to be implemented in VLSI technology.

NORMAL MODE RESPONSE OF INTEGRATING ADC




Figure 4.12

SECTION V

DACs FOR DSP APPLICATIONS



DACs FOR DSP APPLICATIONS

- **DAC ARCHITECTURES**
- **GLITCH REDUCTION BY SEGMENTATION**
- **GLITCH REDUCTION BY DIGITAL OFFSET**
- **DEGLITCHING DACs WITH TRACK-AND-HOLDS**
- **MULTIPLYING DACs**
- **LOGDACs**



SECTION V

DACs FOR DSP APPLICATIONS

In order to be suitable for today's mixed-signal DSP applications, DACs must meet stringent performance specifications. In addition to traditional dc and ac specifications such as resolution, linearity, monotonicity, and settling time, DACs must be evaluated in terms of their frequency-domain dynamic performance such as SNR, THD, and oversampling ratio. DACs must be easy to interface to a variety of DSP chips with little or no additional "glue logic." Communications applications such as modems and digital mobile radio have been a driving force in the development of high performance ac-specified voiceband 12 and 14 bit DACs. The rapidly developing field of digital audio

(especially compact disk players) has been a driving force in the development of low-cost, high-performance DACs having resolutions of 16 to 20 bits with complete SNR and THD dynamic specifications. In the high speed area, the raster-scan display market has a requirement for 8 and 10 bit videodacs complete with on-board memory and other functions which are specific to these applications. In the rapidly evolving direct digital synthesis market (DDS), 12 bit DACs with update rates exceeding 100MSPS are required for the generation of spectrally pure sinewaves. In addition to the above applications, the process control area continues to expand and demand high performance DACs.

5

HIGH-PERFORMANCE DAC APPLICATIONS

- Voiceband: Modems, Speech Synthesis, Digital Mobile Radio — 12 and 14 bits
- Audio: Professional, Compact Disk Players — 16 to 20 bits
- Raster Scan Display Systems — 8 and 10 bits
- Direct Digital Synthesis — 10 and 12 bits, > 100MSPS

Figure 5.1

DAC ARCHITECTURE

Most DACs consist of two basic elements: a set of current or voltage switches, and a method for binary division. Figure 5.2 shows two methods for constructing a DAC. In the first method, the current switches carry equal currents, and the binary scaling is done with an R-2R resistor network. The second method utilizes binarily weighted current sources whose switched outputs are summed together. In the second method, the binary current scaling is often done with an R-2R network in the emitters of the switching transistors.

In DACs fabricated on bipolar processes, current switching is most often done using non-saturating differential transistor pairs as shown in Figure 5.3. Although high performance NPN transistors are more common in bipolar processes, recent complementary bipolar (CB) processes have allowed the development of high speed DACs (such as the AD568 and AD668) based around PNP differential pairs as the basic switch element.

BIPOLAR DAC ARCHITECTURES

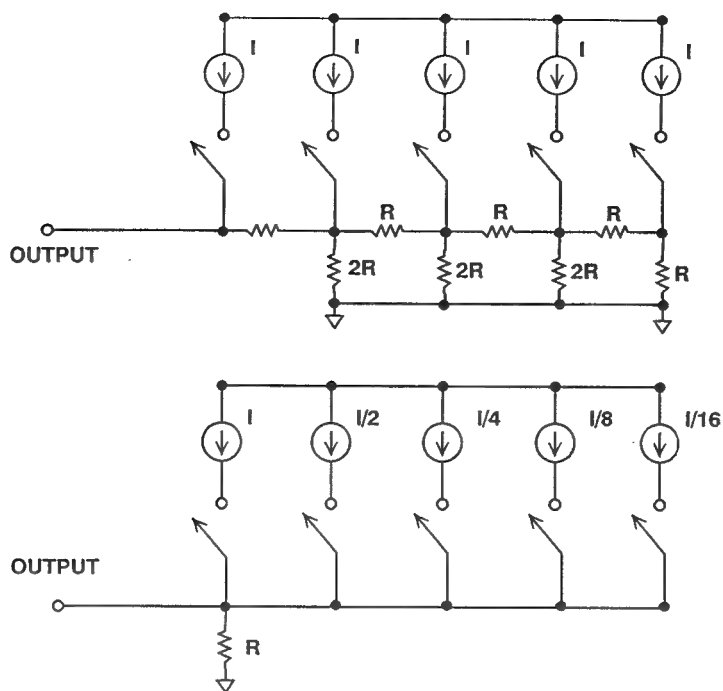


Figure 5.2

TTL AND ECL BIT SWITCHES

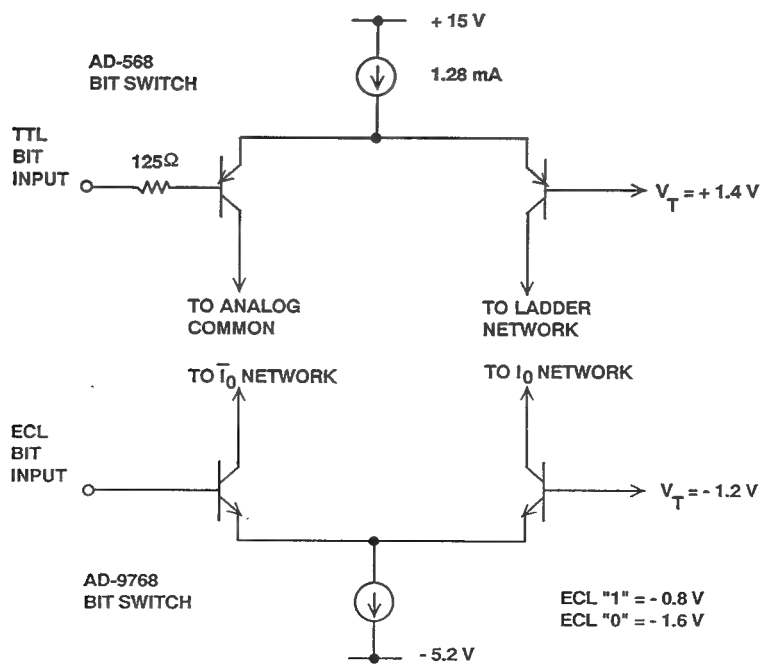


Figure 5.3

CURRENT-STEERING CMOS DAC




Figure 5.4

5

VOLTAGE-SWITCHING CMOS DAC




Figure 5.5

CMOS DACs are most commonly based on the current-mode steering circuit shown in Figure 5.4. An external reference is applied to the V_{ref} pin, and the R-2R ladder divides the input current I into binary-weighted currents as shown. The output drives the virtual ground of an inverting op amp. The finite "on" resistance of the FET switches is compensated for by placing an equivalent compensating FET in series with the feedback resistor R . CMOS DACs can also be realized which operate in the voltage-mode as shown in Figure 5.5.

For high-performance DACs in the voice-band and audio range, BiMOS processes

(bipolar and CMOS devices on the same process) offer the advantages of low-power CMOS for the digital circuits (such as parallel-to-serial converters and latches) along with the low-glitch fast-switching performance of bipolar transistors. A typical current switch cell for such a DAC (the AD1860 18-bit audio DAC) is shown in Figure 5.6. The outputs of CMOS latches are level shifted by the two FETs and converted into a low-level $\pm 0.8V$ differential drive for the NPN differential pair.

BiMOS CURRENT SWITCH




Figure 5.6

GLITCH REDUCTION BY SEGMENTATION

If real estate, cost, power, and capacitance were of no consideration, the ideal “glitch-less” DAC would consist of $2^N - 1$ equally weighted current switches preceded by latches and decoding logic as shown in Figure 5.7. The glitch produced by switching between levels is code-independent and, therefore, does not generate harmonics of the sinewave being reconstructed. A set of latches is required after the binary decoding logic in order to equalize the delays to the actual switches themselves.

The scheme shown in Figure 5.7 is obviously not practical for high resolution DACs, but a significant amount of glitch reduction can be achieved by applying the concept to the first few MSBs. A block diagram of the segmentation technique used in the AD1860 18 bit audio DAC is shown in Figure 5.8.

The AD1860 uses a combination of segmentation and R-2R division to achieve excellent linearity and low distortion. The four MSBs are decoded into a 15 bit “thermometer” code after the serial-to-parallel conversion. The 15 decoded lines (representing the 4 MSBs) and the 14 LSB lines are then fed to a 29 bit latch. The 15 thermometer decoded lines each drive current switches having equal weights. The 14 LSB lines drive a conventional binary-weighted R-2R DAC. This combination of segmentation and conventional R-2R architecture along with laser trimmed thin film resistors allows the AD1860 to meet stringent audio specifications without the need for an external SHA deglitcher. A summary of key performance specifications for the AD1860 is given in Figure 5.9.

IDEAL DAC FOR MINIMUM GLITCH




Figure 5.7

SEGMENTATION IN AD1861 18-BIT AUDIO DAC

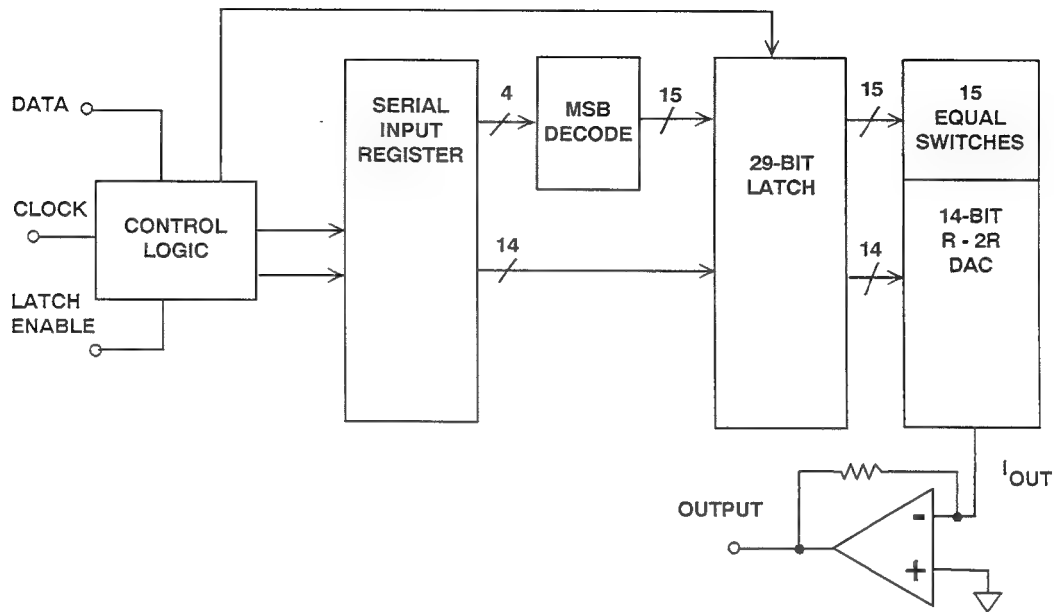


Figure 5.8

AD1861 18-BIT AUDIO DAC KEY SPECIFICATIONS

- 108dB SNR
- 0.002% THD + N @ 0dB Signal Amplitude
- Up to 16x Oversampling Capability (768kSPS)
- $\pm 3V$ or $\pm 1mA$ Output Capability
- 110mW Power Dissipation
- 16 Pin DIP Package

Figure 5.9

Figure 5.10 shows how the segmentation architecture is applied to the AD7840 14 bit CMOS DAC. The three MSBs of the input binary word are decoded into a 7 bit thermometer code to drive the seven switches

A-G. The 11 LSBs switch an 11 bit R-2R ladder structure. A functional block diagram of the AD7840 is shown in Figure 5.11, and key specifications are given in Figure 5.12.

SEGMENTATION IN THE AD7840 14-BIT DAC

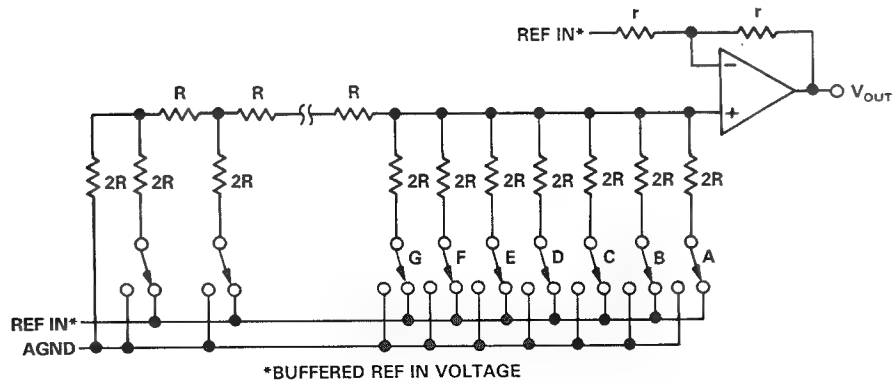


Figure 5.10

AD7840 FUNCTIONAL BLOCK DIAGRAM

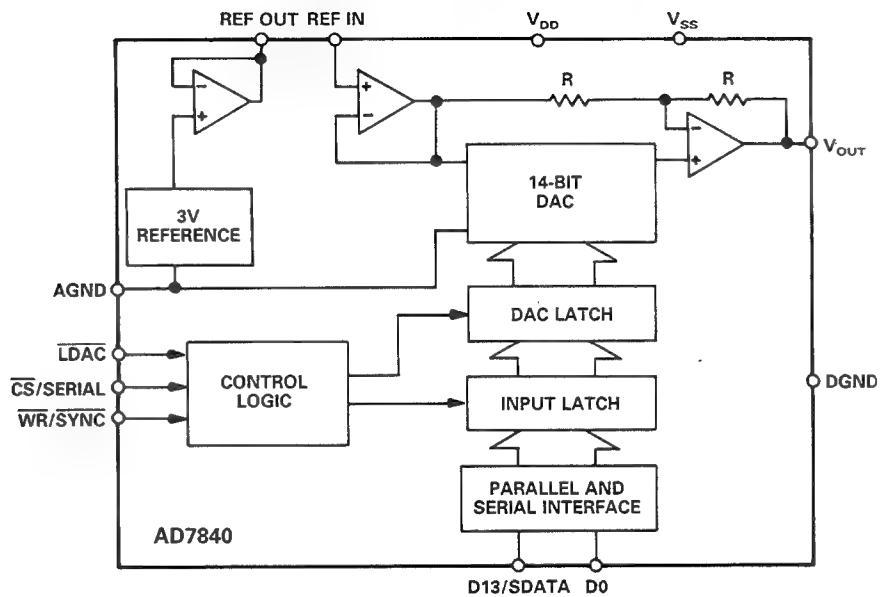


Figure 5.11

AD7840 14 BIT CMOS DAC KEY SPECIFICATIONS

- Complete 14 Bit Voltage Output DAC
- Parallel and Serial Interface Capability
- 80dB SNR
- Easy Interface to DSP Processors
- 100kSPS Update Rate
- 100mW, $\pm 5V$ Supplies

Figure 5.12

Segmentation can also be applied to voltage-mode output Bi-CMOS DACs as shown in Figure 5.13 for the AD569 16 bit DAC. The DAC consists of two resistor strings, each of which is divided into 256 equal segments. The 8 MSBs of the digital input word select one of the 256 segments on the first string. The taps at the top and bottom of the selected segment are connected to the inputs of the two buffer amplifiers A1 and A2. The buffered voltages from the segment endpoints are applied across the second resistor string, where the 8 LSBs of the digital input word select one of the 256 output taps. Output buffer A3 buffers this

voltage and delivers it to the output. Buffer amplifiers A1 and A2 leap-frog up the first string to preserve monotonicity at the segment boundaries. For example, when increasing the digital code from $00FF_H$ to 0100_H , (the first segment boundary), A1 remains connected to the same tap on the first string, while A2 jumps over it and is connected to the tap which becomes the top of the next segment. This design guarantees monotonicity even if the amplifiers have offset voltages. In fact, amplifier offset only contributes to integral linearity error. Key specifications for the AD569 are summarized in Figure 5.14.

SEGMENTATION IN THE AD569 16-BIT DAC

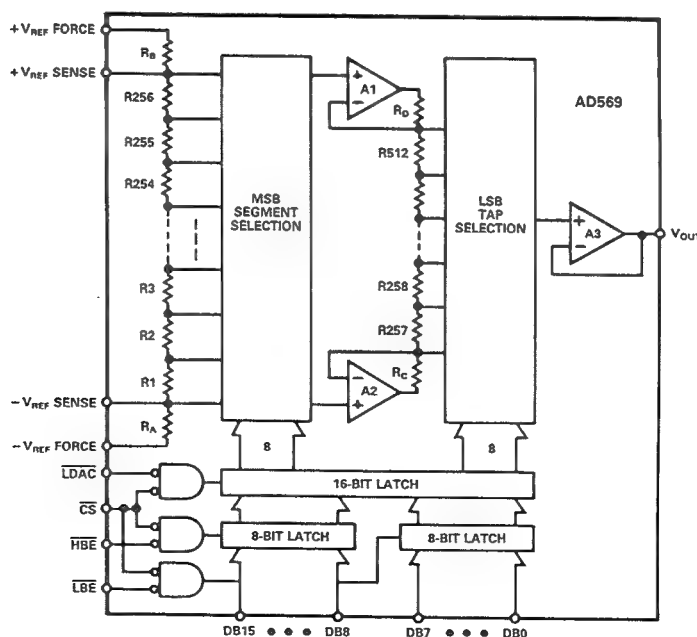


Figure 5.13

AD569 DAC KEY SPECIFICATIONS

- Guaranteed 16 Bit Monotonicity
- Linear to 13 Bits
- 3 μ s Settling to 16 Bits (0.001%)
- 500nV-sec Glitch Impulse Area
- 8 and 16 Bit Bus Compatibility
- Low Power: 250mW

Figure 5.14

The AD7846 is another 16 bit DAC which uses the segmented architecture, and a block diagram is shown in Figure 5.15. The 4 MSBs in the DAC latch select one of the segments in a 16-resistor string. Both taps of the segment are buffered by amplifiers and

fed to a 12 bit DAC which provides a further 12 bits of resolution. As with the AD569, a leap-frog approach is used to prevent non-monotonicity at the segment switching points. A summary of key specifications for the AD7846 is given in Figure 5.16.

SEGMENTATION IN THE AD7846 16-BIT DAC

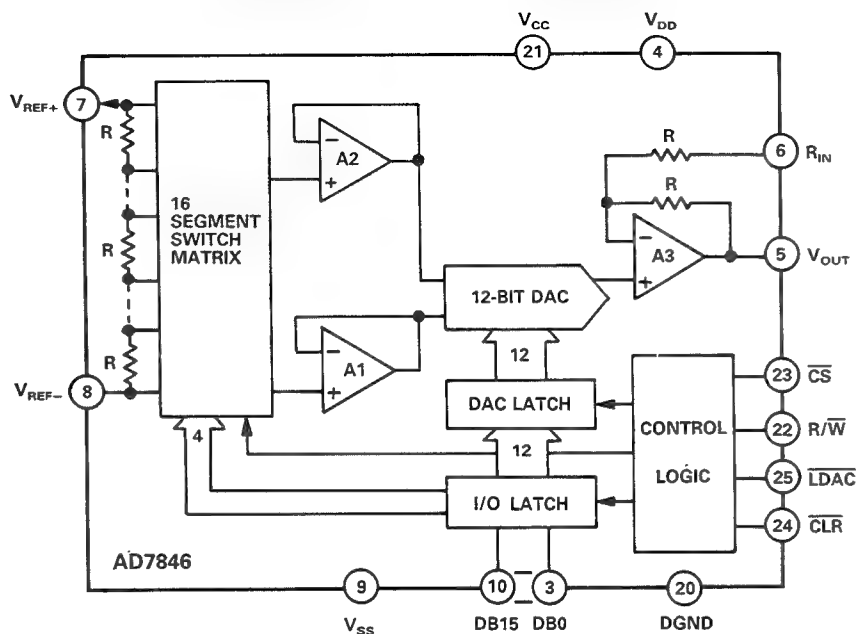


Figure 5.15

AD 7846 DAC KEY SPECIFICATIONS

- Guaranteed 16 Bit Monotonicity
- Linear to 15 Bits
- 9 μ s Settling to 0.003%
- 400nV-sec Glitch Impulse Area
- Microprocessor Compatible With Readback Capability
- Low Power: 100mW

Figure 5.16

GLITCH REDUCTION BY DIGITAL OFFSET

Regardless of the architecture used to design a high performance DAC, the most troublesome code-dependent glitch typically occurs at the midscale code transition, i.e. from 0111...1 to 1000...0. In an audio system which operates with bipolar signals, the midscale glitch noise is particularly troublesome, since it can introduce distortion for

very low-level passages. If a small digital offset is added to the DAC input, then the DAC midscale glitch noise will occur at a slightly higher input signal which is less objectionable. Unfortunately, one end of the DACs range will be clipped by an amount equal to the injected digital offset as shown in Figure 5.17.

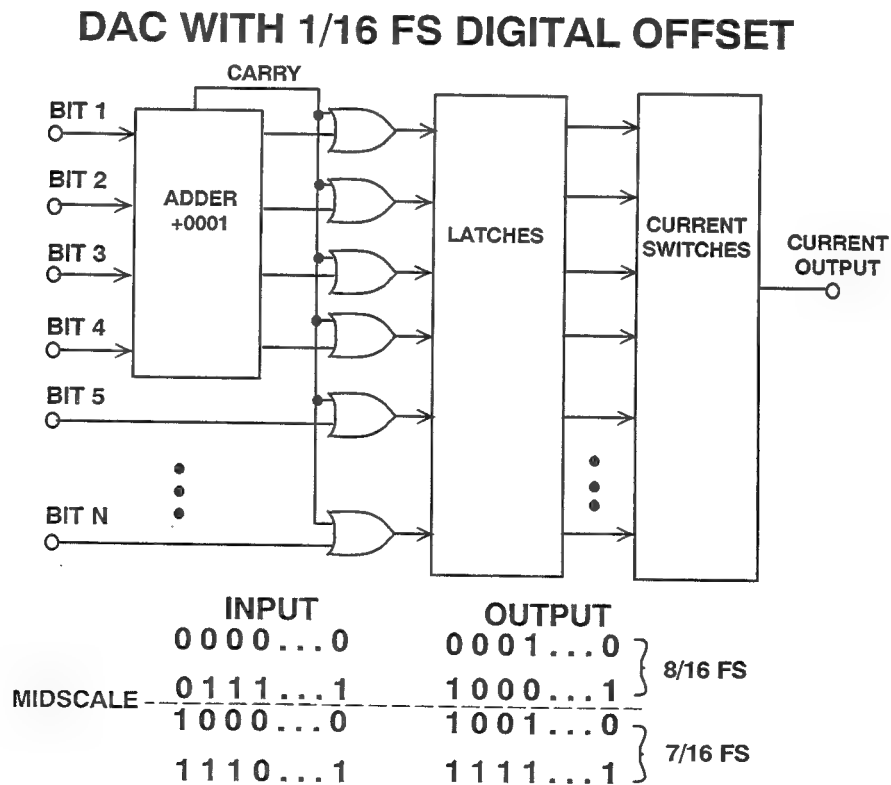


Figure 5.17

The AD1862 20 bit digital audio DAC uses a combination of segmentation and digital offset to achieve a high level of performance. The novel architecture prevents clipping and allows the full range of the 20 bit DAC to be utilized. A block diagram of the AD1862 is shown in Figure 5.18. The digital offset is accomplished by adding 0001 (1/16th full-scale) to the four MSBs. The three MSBs are then segmented into a 7 bit thermometer code output which is latched and then drives seven equal current switches. Bit 4 (after the addition), and bits 5 through 20 are

latched and then drive a conventional R-2R DAC. In order to prevent clipping at the positive end of the range, the carry output of the adder drives an additional current switch having a weight corresponding to bit 4. Finally, an offset current equal to 1/16th fullscale is subtracted from the DAC output to compensate for the constant digital offset. This architecture results in exceptional THD + N performance as shown in 5.19. Key performance specifications for the AD1862 are summarized in Figure 5.20.

AD1862 20-BIT AUDIO DAC DIGITAL OFFSET AND SEGMENTATION




Figure 5.18

THD + N VERSUS INPUT FREQUENCY FOR AD1862 20-BIT AUDIO DAC




Figure 5.19

AD1862 20-BIT AUDIO DAC KEY SPECIFICATIONS

- 119dB SNR
- 0.0016% THD + N @ 0dB Signal Amplitude
- 16x Oversampling Capability (705.6kSPS)
- ± 1 dB Gain Linearity @ -90dB Amplitude
- ± 1 mA Output Current
- 288mW Power Dissipation
- 16 Pin DIP Package

Figure 5.20

DEGLITCHING DACs WITH TRACK-AND HOLDS

As has been previously mentioned, code-dependent glitches can be effectively removed with a SHA as shown in Figure 5.21. Just prior to latching new data into the DAC, the SHA is placed in the hold mode so that the DAC switching glitches are isolated from the output. The switching transients produced by the SHA are code-independent and occur at the update frequency, hence, they

are easily filterable. Although the terms *sample-and-hold* and *track-and-hold* are often used interchangeably, some sample-and-holds have poor track- or sample-mode performance. In order to be used as a DAC deglitcher, the device must function as a true track-and-hold with good track-mode performance.




Figure 5.21

A high-performance low-cost SHA deglitcher circuit is shown in Figure 5.22. This circuit will operate with high-speed DACs such as the AD568 and AD668 (12 bit DACs) at update rates up to 10Mhz with harmonic suppression of 70 to 75dB.

A SHA deglitcher suitable for 14 bit performance at 100kSPS is shown connected to the AD7840 14 bit DAC in Figure 5.23. Input frequencies up to 20kHz can be reconstructed while maintaining an SNR of 82dB.

10 MSPS DEGLITCHING CIRCUIT

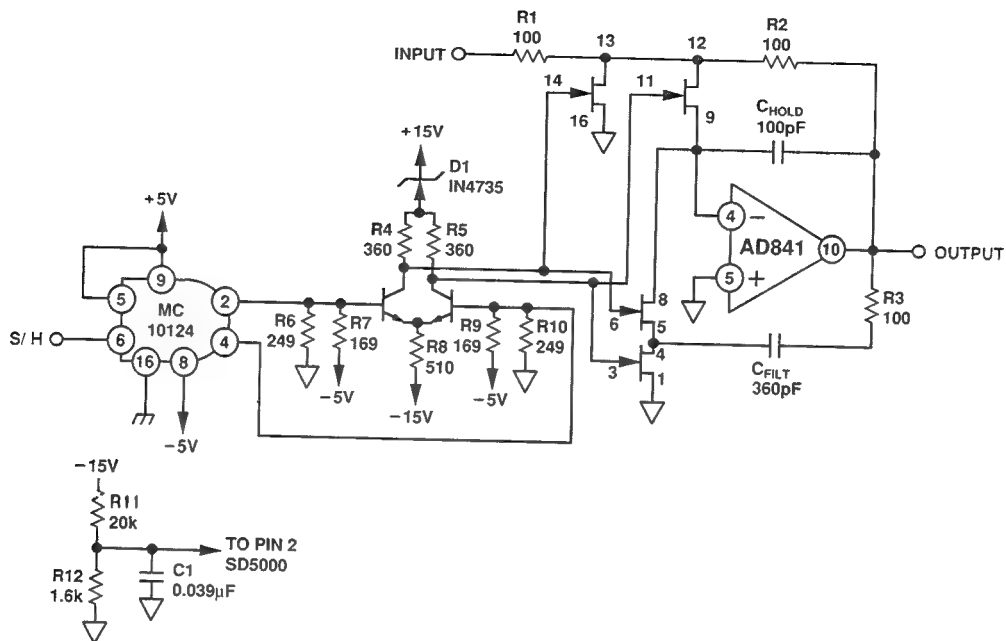


Figure 5.22

DEGLITCHING A 14-BIT, 100kSPS DAC

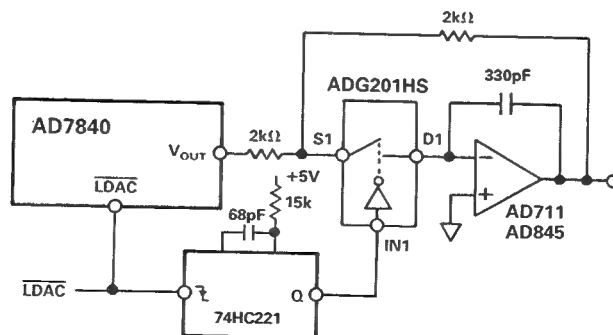


Figure 5.23

SUMMARY OF GLITCH REDUCTION TECHNIQUES

- Low Voltage Swings to Drive Switches
- Current-Mode Switching
- Input Latches Directly Before Switches to Reduce Skew
- Segmentation of Higher-Order Bits
- Digital Offset
- SHA Deglitching

Figure 5.24

MULTIPLYING DACs

Virtually all DACs produce an output voltage (or current) which is proportional to the product of a reference voltage (either internal or external to the device) and the digital input word. A multiplying DAC is simply a DAC whose reference voltage can be varied externally over a specified range (see Figure 5.25). This feature can be used in several ways. The DAC can be used as a digital potentiometer (or programmable gain amplifier) which attenuates the signal applied to the reference input terminal proportionally to the digital input word. The ability to digitally control the amplitude of ac signals depends on the bandwidth of the reference voltage input. In order to be useful as a general purpose digital attenuator, the DAC

should be able to operate with a positive, negative, or zero reference voltage. The CMOS DAC architecture of the AD7845 shown in Figure 5.26 allows this flexibility, while some DACs fabricated on bipolar processes typically allow reference voltages of one polarity only, and sometimes values near zero are restricted.

True MDACs constructed on CMOS processes with R-2R ladder networks allow for bipolar reference inputs, and if the DAC is also operated in the bipolar output mode, then true four-quadrant multiplication can be realized as shown in Figure 5.27. Key features of the AD7845 are summarized in Figure 5.28.

5

MULTIPLYING DAC




Figure 5.25

AD7845 12-BIT MULTIPLYING DAC

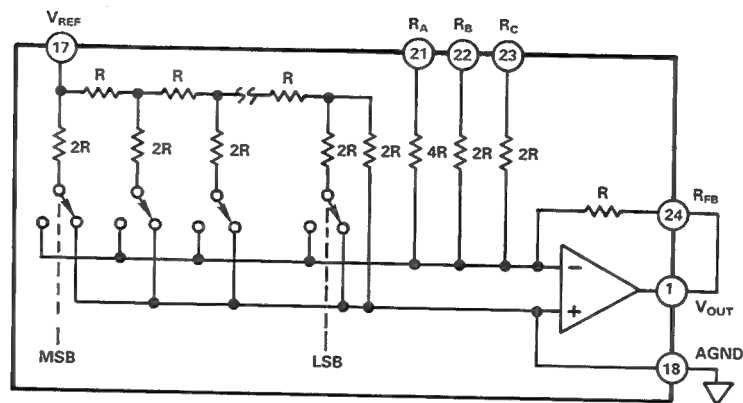


Figure 5.26

4-QUADRANT MDAC CONFIGURATION

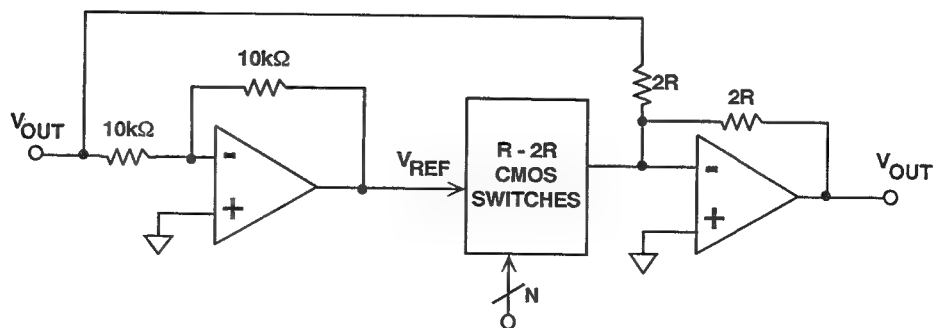


Figure 5.27

KEY FEATURES OF THE AD7845 12-BIT MDAC

- 12 Bit CMOS MDAC with Output Amplifier
- 4-Quadrant Multiplication
- 250kHz Full Power Reference Voltage Bandwidth (20V p-p)
- 90dB Total Harmonic Distortion for 1kHz 6V rms Sinewave

Figure 5.28

LOGDACs

A LOGDAC is a multiplying DAC with a gain proportional to the exponential of the digital input. Equal changes of digital input produce equal *ratios* of analog gain change. In other words, the weight of the least significant bit (LSB) is expressed in dB relative to fullscale. A linear DAC can be used to generate a logarithmic attenuation function as shown in Figure 5.29 for a 6-bit DAC. This is accomplished by properly selecting only 7 out of the 64 possible input codes as shown. The proper 6-bit code can thus be specified with a 3-bit code, and a ROM decoder can be used to make the translation. For this example, the weight of the LSB corresponds to 6dB.

A block diagram of the AD7111 8-bit LOGDAC is shown in Figure 5.30. This DAC is a CMOS multiplying DAC which can attenuate an analog input signal over the range 0 to 88.5dB in 0.375dB steps. The circuit consists of a 17-bit R-2R CMOS multiplying DAC with decoding logic which translates the 8-bit binary input word into a 17-bit word which actually drives the internal DAC. An input code of 0 (decimal) corresponds to 0dB attenuation, while an input code of 239 (decimal) corresponds to an attenuation of 88.5dB. For input codes of 240 through 255, the output is zero. Key specifications for the AD7111 LOGDAC are shown in Figure 5.31.

5


ATTENUATION VERSUS SELECTED INPUT CODES
FOR LINEAR 6-BIT MDAC

Figure 5.29

AD7111 8-BIT LOGDAC

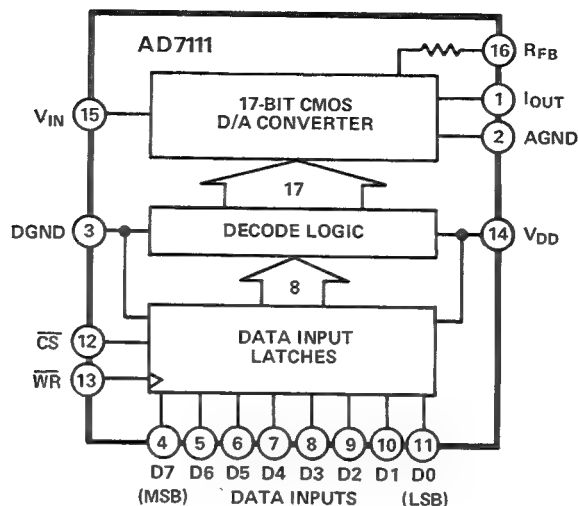


Figure 5.30

KEY SPECIFICATIONS FOR THE AD7111 LOGDAC

- 88.5dB Dynamic Range
- 0.375dB Resolution, 8-Bit Input, 239 Levels
- 90dB THD for 6V rms 1kHz Input, 0dB Attenuation

Figure 5.31

APPLICATIONS OF MDACs AND LOGDACs

- Digital Attenuators
- Programmable Power Supplies
- Programmable Gain Amplifiers
- Digitally Controlled AGC Systems

Figure 5.32

REFERENCES

1. **High Speed Design Seminar**, Analog Devices, 1990.
2. Phil Burton, **CMOS DAC Application Guide**, 3rd Edition, Analog Devices, 1984.
3. Daniel H. Sheingold, Editor, **Analog-Digital Conversion Handbook**, Prentice-Hall, 1986.



SECTION VI

SIGMA-DELTA ADCs AND DACs

SIGMA-DELTA ADCs AND DACs

- **SIGMA-DELTA OVERVIEW**
- **OVERSAMPLING**
- **SIGMA-DELTA MODULATORS AND QUANTIZATION NOISE SHAPING**
- **DIGITAL FILTERING AND DECIMATION**
- **IDLING PATTERNS AND TONAL CONSIDERATIONS FOR SIGMA-DELTA ADCs**
- **HIGHER ORDER MODULATOR LOOPS**
- **DESCRIPTION OF AD1879 18 BIT SIGMA-DELTA AUDIO ADC**
- **SIGMA-DELTA ADCs FOR LOW FREQUENCY MEASUREMENT APPLICATIONS**
- **SIGMA-DELTA DACs**
- **THE ADSP-28msp02 SIGMA-DELTA CODEC**
- **MULTISTAGE NOISE SHAPING (MASH) SIGMA-DELTA CONVERTERS**
- **MULTI-BIT SIGMA-DELTA CONVERTERS**
- **SIGMA-DELTA SUMMARY**

SECTION VI

SIGMA-DELTA ADCs AND DACs

SIGMA-DELTA OVERVIEW

Within the last several years, the sigma-delta architecture has become more and more popular for realizing high-resolution ADCs in mixed-signal VLSI processes. Until recently, however, the process technology needed to make these devices commercially viable has not been available. Now that 1 micron and smaller CMOS geometries are manufacturable, sigma-delta converters will become even more prolific in certain types of applications, especially mixed-signal ICs which combine the ADC, DAC, and DSP functions on a single chip.

Conceptually, the sigma-delta architecture is more digital than analog intensive. This does not, however, minimize the importance of the analog portion of the sigma-delta ADC. The design of a fifth-order sigma-delta modulator (as in the AD1879 dual 18 bit ADC) is certainly not a trivial matter, and neither is the digital filter. The sigma-delta converter

is inherently an oversampling converter, although oversampling is just one of the techniques contributing to the overall performance. Basically, a sigma-delta converter digitizes an analog signal with a very low resolution (1 bit) ADC at a very high sampling rate. By using oversampling techniques in conjunction with noise shaping and digital filtering, the effective resolution is increased. Decimation is then used to reduce the effective sampling rate at the ADC output. The sigma-delta ADC exhibits excellent differential and integral linearity due to the linearity of the 1 bit quantizer and DAC, and no trimming is required as in other ADC architectures.

The key concepts involved in understanding the operation of sigma-delta converters are oversampling, noise shaping (using a sigma-delta modulator), digital filtering, and decimation.

6

SIGMA-DELTA CONCEPTS

- Ideal Topology for Mixed Signal VLSI Chips
- Oversampling
- Noise-Spectrum Shaping Using Sigma-Delta Modulator
- Digital Filtering
- Decimation
- 16 Bits and Higher Resolution Possible

Figure 6.1

OVERSAMPLING

The concept of oversampling has been previously discussed in Section III, and is illustrated again in Figure 6.2 and 6.3. As was discussed, one significant benefit of oversampling is that the rolloff requirements on the analog antialiasing filter are relaxed. The quantization noise (rms value over Nyquist bandwidth is $q/\sqrt{12}$, where q is the weight of the LSB) which falls between $f_s/2$ and $kf_s/2$ is removed from the output by the digital filter (k is the oversampling ratio). This has the effect of increasing the overall signal-to-noise ratio by an amount equal to $10\log_{10}(k)$. Unfortunately this is a high price to pay for extra resolution, as an oversampling ratio of

4 is required just to increase the signal-to-noise ratio by a modest 6dB (1 bit). To keep the oversampling ratio within reasonable bounds, it is possible to shape the frequency spectrum of the quantization noise so that the majority of the noise lies between $f_s/2$ and $kf_s/2$, and only a small portion is left between dc and $f_s/2$. This is precisely what a sigma-delta modulator does in a sigma-delta ADC. After the noise spectrum is shaped by the modulator, the digital filter can then remove the bulk of the quantization noise energy, and the overall signal-to-noise ratio (hence the dynamic range) is dramatically increased.

NYQUIST SAMPLING WITH ANALOG LOWPASS FILTER




Figure 6.2

OVERSAMPLING WITH ANALOG AND DIGITAL FILTERING




Figure 6.3

SIGMA-DELTA MODULATORS AND QUANTIZATION NOISE SHAPING

A block diagram of a first-order sigma-delta ADC is shown in Figure 6.4. The first part of the converter is the sigma-delta modulator which converts the input signal into a continuous serial stream of 1's and 0's at a rate determined by the sampling clock frequency, Kf_s . The 1-bit DAC is driven by the serial output data stream, and the DAC output is subtracted from the input signal. Feedback control theory tells us that the average value of the DAC output (hence the serial bit stream) must approach that of the input signal if the loop has enough gain. The integrator can be represented in the frequency domain by a filter whose amplitude response is proportional to $1/f$, where f is the input frequency. Since the chopper-like action of the clocked, latched comparator converts the input signal to a high-frequency ac signal, varying about the average value of the input, the effective quantization noise at

low frequencies is greatly reduced (the integrator looks like a high-pass filter to quantization noise). The exact frequency spectrum of the resulting noise depends on the sampling rate, the integrator time constant, and the precise span of the voltage fed back.

For any given input value in a single sampling interval, the data from the 1-bit ADC is virtually meaningless. Only when a large number of samples are averaged, will a meaningful value result. The sigma-delta modulator is very difficult to analyze in the time domain because of this apparent randomness of the single-bit data output. If the input signal is near positive fullscale, it is clear that there will be more 1's than 0's in the bit stream. Likewise, for signals near negative fullscale, there will be more 0's than 1's in the bit stream. For signals near midscale, there will be approximately an equal number of 1's and 0's. Figure 6.5

FIRST-ORDER SIGMA-DELTA ADC

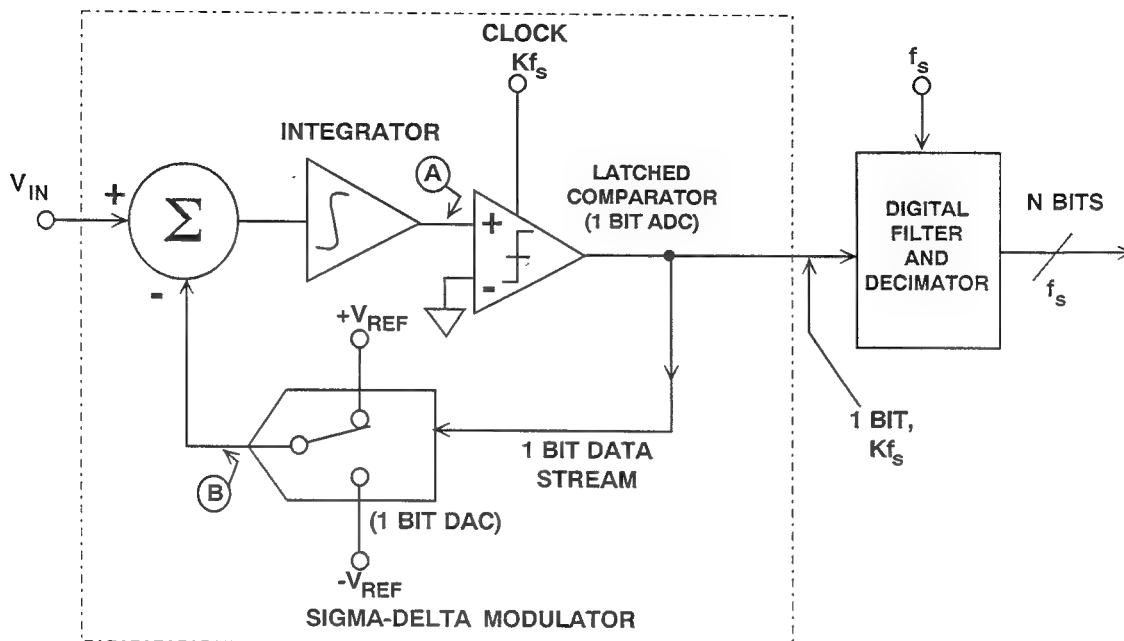


Figure 6.4

SIGMA-DELTA MODULATOR WAVEFORMS

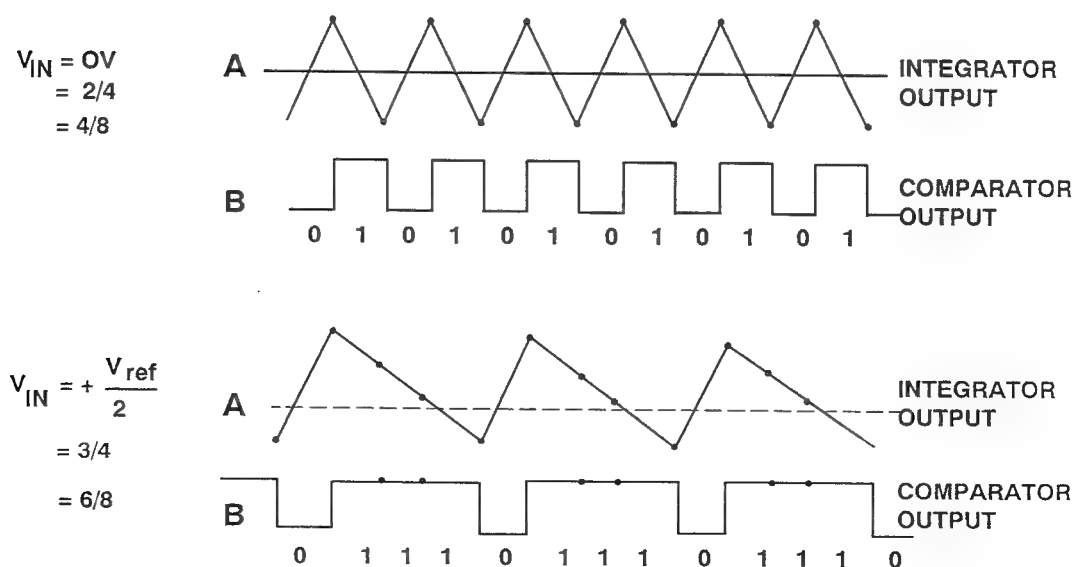


Figure 6.5

shows the output of the integrator for two input conditions. The first is for an input of zero (midscale). To decode the output, pass the output samples through a simple digital lowpass filter that averages every four samples. The output of the filter is $2/4$. This value represents bipolar zero. If more samples are averaged, more dynamic range is achieved. For example, averaging 4 samples gives 2 bits of resolution, while averaging 8 samples yields $4/8$, or 3 bits of resolution. In the bottom waveform of Figure 6.5, the average obtained for 4 samples is $3/4$, and the average for 8 samples is $6/8$.

The sigma-delta ADC can also be viewed as a synchronous voltage-to-frequency converter followed by a counter. If the number of 1's in the output data stream is counted over a sufficient number of samples, the counter output will represent the digital value of the input. Obviously, this method of averaging will only work for dc or very slowly changing input signals. In addition, 2^N clock cycles must be counted in order to achieve N-bit effective resolution, thereby severely limiting the effective sampling rate.

Further analysis of the sigma-delta archi-

ture is best done in the frequency domain using the linear model shown in Figure 6.6. Note that the integrator is represented as an analog filter with a given transfer function $H(f)$. The transfer function has an amplitude response which is inversely proportional to the input frequency. The quantizer is modeled as a gain stage followed by the addition of quantization noise. One of the advantages of using frequency domain analysis is that algebra can be used to describe the signals. The output value y can be represented as the difference $x - y$ from the summing node at the input multiplied by the transfer function of the analog filter (integrator), multiplied by the gain block, and then added with the quantization noise Q . If we set the gain to 1, and the transfer function is represented as $1/f$, the following mathematical relationship results:

$$y = \frac{x - y}{f} + Q, \text{ or by rearranging,}$$

$$y = \frac{x}{f + 1} + \frac{Qf}{f + 1}.$$

6

FREQUENCY DOMAIN LINEARIZED MODEL OF A SIGMA DELTA MODULATOR

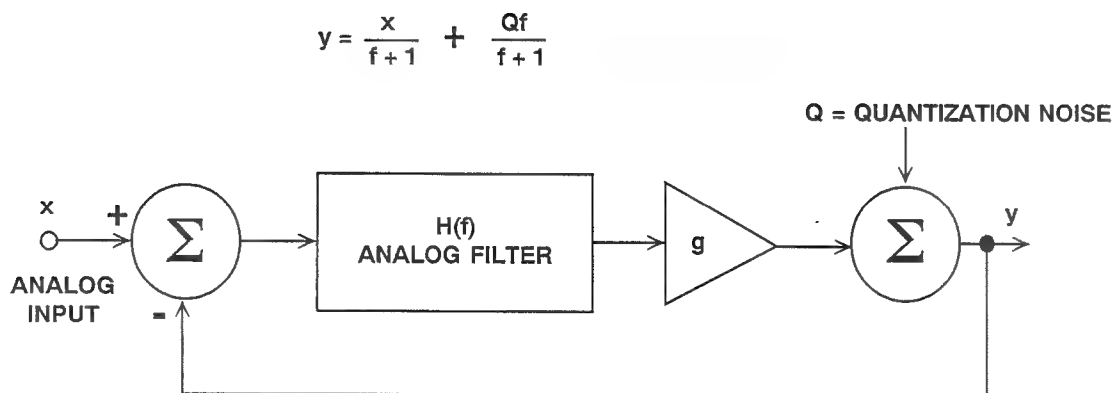


Figure 6.6

Note that as frequency f approaches 0, the output approaches x with no noise component. At higher frequencies, the value of x is reduced, and the value of the noise component is increased. For high frequency inputs, the output consists primarily of quantization noise. In essence, the analog filter has a low pass effect on the signal and a high pass effect on the noise component. For this reason, the analog filter of the modulator can be viewed as a noise shaping filter as shown in Figure 6.7.

As with analog filters in general, higher order filters offer better performance. This is also true of the sigma-delta modulator, provided certain precautions are taken. A second order sigma-delta modulator is shown in Figure 6.8, and a comparison between the noise shaping functions is shown in Figure 6.9. Figure 6.10 shows a plot of the corresponding in-band signal-to-noise ratio (dynamic range) as a function of the oversampling ratio for a first and second order modulator. Note that the first order transfer

function has a slope of 9dB per octave, while the second order transfer function slope is 15dB per octave. Higher order modulators (greater than second order) can realize even better performance, but the simple linear model must be used with great care, and sophisticated design techniques are required in order to insure stability. The curve shown in Figure 6.10 for the third-order loop represents an unrealizable condition and is shown for reference only.

The curves in Figure 6.10 can be used to determine the approximate ADC resolution achievable, given the modulator order and the oversampling rate. For instance, if the oversampling rate is $64x$, an ideal second order system is capable of providing a signal to noise ratio of about 80dB. This implies an ADC resolution of approximately 13 bits. Although the filtering done by the digital filter can be done to any degree of precision desirable, it would be pointless to carry more than 13 binary bits to the outside world. Additional bits would carry no useful signal information, and would be buried in noise.

SHAPED QUANTIZATION NOISE DISTRIBUTION




Figure 6.7

SECOND-ORDER SIGMA-DELTA ADC

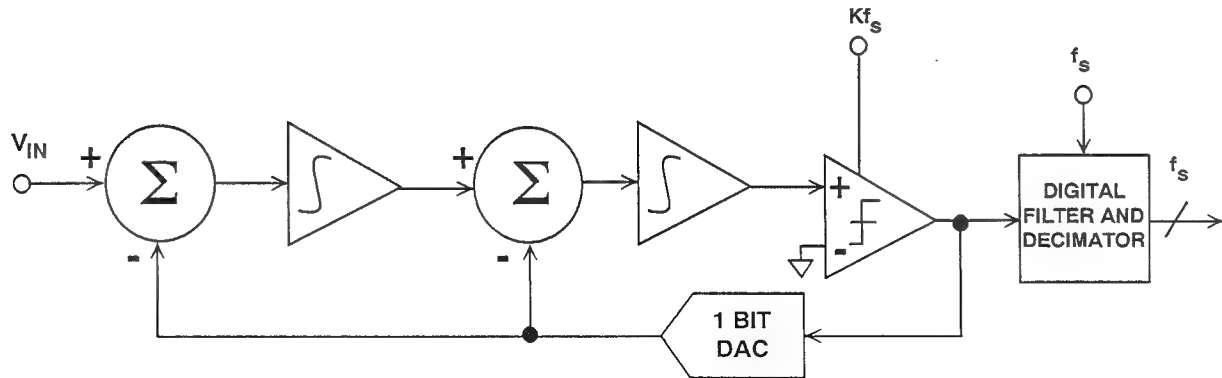


Figure 6.8

6

FIRST AND SECOND-ORDER NOISE SHAPING FUNCTIONS




Figure 6.9

SNR VERSUS OVERSAMPLING RATIO FOR FIRST, SECOND, AND THIRD-ORDER LOOPS




Figure 6.10

DIGITAL FILTERING AND DECIMATION

After the quantization noise has been shaped by the modulator and pushed into the frequencies above the band of interest, digital filtering techniques can be applied to this shaped quantization noise as shown in Figure 6.11. The purpose of the digital filter is twofold. First, it must act as an anti-aliasing filter with respect to the final sampling rate, f_s . Second, it must filter out the higher frequency noise produced by the noise-shaping process of the sigma-delta modulator.

The final data rate reduction is performed by digitally resampling the filtered output using a process called decimation. The decimation of a discrete-time signal is shown in Figure 6.12, where the sampling rate of the input signal $x(n)$ is at a rate which is to be reduced by a factor of 4. The signal is resampled at the lower rate (the decimation rate), $s(n)$. Decimation can also be viewed as the method by which the redundant signal

information introduced by the oversampling process is removed.

In sigma-delta ADCs it is quite common to combine the decimation function with the digital filtering function. This results in an increase in computational efficiency if done correctly.

Recall that a finite impulse filter (FIR) simply computes a moving weighted average (the weighting being determined by the individual filter coefficients) of the input samples. Normally, there is one filter output for every input sample. If, however, we wish to decimate the filter output by digitally resampling at a lower rate, it is no longer necessary to compute a filter output for every input sample. Instead, we only compute filter outputs at the lower decimation rate, thereby achieving considerable efficiency in the computational process.

If, however, an infinite impulse response (IIR) filter is used, it is necessary to compute

EFFECTS OF DIGITAL FILTERING ON SHAPED QUANTIZATION NOISE




Figure 6.11

6

DECIMATION OF A DISCRETE-TIME SIGNAL




Figure 6.12

an output for every input (because of the feedback term), and therefore the decimation cannot be performed as part of the digital filtering process. In some sigma-delta ADC designs, the filtering is performed in two stages. If both FIR and IIR filters are used, the decimation is performed in the first FIR stage, and the final filtering is done in the final IIR stage. If FIR filters are used for both stages, it is usually more efficient to split the decimation between the two filter stages.

From the above discussion it should be clear that the design of a sigma-delta ADC digital filter involves many tradeoffs. FIR filters lend themselves to decimation, are always stable, and have linear phase characteristics (extremely important in audio and

some telemetry applications). Although they are typically easier to design, they usually require more stages to realize a given transfer characteristic than a corresponding IIR filter. On the other hand, the IIR filter employs feedback which eliminates the possibility of decimation within the filter, but makes the filter more efficient (better filter performance with fewer calculations). The feedback used in IIR filters can lead to a potentially unstable filter implementation. Also, the IIR filter (which will closely emulate filter functions realized in the analog domain) exhibits non-linear phase characteristics. Because of the stability issues and the quantization effects in the feedback loop, IIR filters are more complicated to design correctly.

SIGMA-DELTA ADC DIGITAL FILTERING AND DECIMATION

FIR Filters:

- Easy to Design
- Easy to Incorporate Decimation
- Linear Phase Response
- Large Number of Coefficients May Be Required

IIR Filters:

- Stability, Overflow Considerations
- Cannot Decimate Internally Due to Feedback
- More Efficient than FIR Filters
- Non-Linear Phase Response

Combinations:

- 2-Stage FIR Filters
- FIR Filter Followed by IIR Filter
- 2-Stage IIR Filters

Figure 6.13

IDLING PATTERN AND TONAL CONSIDERATIONS FOR SIGMA-DELTA ADCs

In our discussion of sigma-delta ADCs up to this point, we have made the assumption that the quantization noise produced by the sigma-delta modulator is random and uncorrelated with the input signal. Unfortunately, this is not entirely the case, especially for the first-order modulator. Consider the case where we are averaging 16 samples of the modulator output in a 4 bit sigma-delta ADC. Figure 6.14 shows the bit pattern for two input signal conditions: an input signal having the value 8/16, and an input signal having the value 9/16. In the case of the

9/16 signal, the modulator output bit pattern has an extra "1" every 16th output. This will produce energy at $f_s/16$, which translates into an unwanted tone. If the oversampling ratio is less than 16, this tone will fall into the passband. Figure 6.15 shows the correlated idling pattern behavior for a first order sigma-delta modulator, and Figure 6.16 shows the relatively uncorrelated pattern for a second-order modulator. For this reason, virtually all sigma-delta ADCs contain at least a second-order modulator loop.

REPETITIVE BIT PATTERN IN SIGMA-DELTA MODULATOR OUTPUT


16 SAMPLES OF SIGMA-DELTA MODULATOR DATA OUTPUT STREAM		BINARY EQUIVALENT
1 0 1 0 1 0 1 0 1 0 1 0 1 0 1 0 ... 8/16	=	1000
1 0 1 0 1 0 1 0 1 0 1 0 1 1 ... 9/16	=	1001
<div style="text-align: center;">  <p>REPEATS EVERY 16 SAMPLES</p> </div>		

Figure 6.14

IDLING PATTERNS FOR FIRST-ORDER SIGMA-DELTA MODULATOR (INTEGRATOR OUTPUT)

IDLE BEHAVIOR WITH 0 VOLTS INPUT



IDLE BEHAVIOR WITH DC INPUT SHOWING CORRELATED IDLING PATTERN




Figure 6.15

IDLING PATTERNS FOR SECOND-ORDER SIGMA-DELTA MODULATOR (SECOND INTEGRATOR OUTPUT)




Figure 6.16

HIGHER ORDER MODULATOR LOOPS

In order to achieve wide dynamic range, sigma-delta modulator loops greater than second-order are necessary, but present real design challenges. First of all, the simple linear models previously discussed are no longer fully accurate. Loops of order greater than two are generally not guaranteed to be stable under all input conditions. The instability arises because the comparator is a non-linear element whose effective “gain” varies inversely with the input level. This mecha-

nism for instability causes the following behavior: if the loop is operating normally, and a large signal is applied to the input that overloads the loop, the average gain of the comparator is reduced. The reduction in comparator gain in the linear model causes loop instability. This causes instability even when the signal that caused it is removed. In actual practice, such a circuit would normally oscillate on power-up due to initial conditions caused by turn-on transients.

HIGHER ORDER LOOP CONSIDERATIONS (>2)

- Increased Dynamic Range and Resolution is Achievable
- Higher Order Loops Minimize Idling Patterns and Tones
- Difficult to Analyze and Stabilize
- Non-Linear Stabilization Techniques Can Be Used Successfully: AD1879 18 Bit, 5th Order ADC

Figure 6.17

Instability in the AD1879 fifth-order modulator is sensed digitally by counting the number of consecutive ones or zeros in the modulator bit stream. A sufficiently long string of either ones or zeros indicates modu-

lator instability. This triggers circuitry which resets the state in the integrators to put the modulator into a stable operating condition.

DESCRIPTION OF AD1879 18 BIT SIGMA-DELTA AUDIO ADC

The AD1879 is a state-of-the-art dual 18 bit sigma-delta ADC designed to meet the stringent requirements of professional digital audio. A block diagram of the device is shown in Figure 6.18, and performance specifications are given in Figure 6.19. The modulator is a fifth-order switched capacitor design which shapes the noise spectrum as

shown in Figure 6.20. The oversampling ratio is 64x, which places the oversampling frequency at 3.072MHz for the standard audio sampling rate of 48kHz. Because of the high oversampling ratio, a single-pole analog antialiasing filter is sufficient at the input of the ADC.

AD1879 DUAL 18-BIT SIGMA-DELTA ADC

6

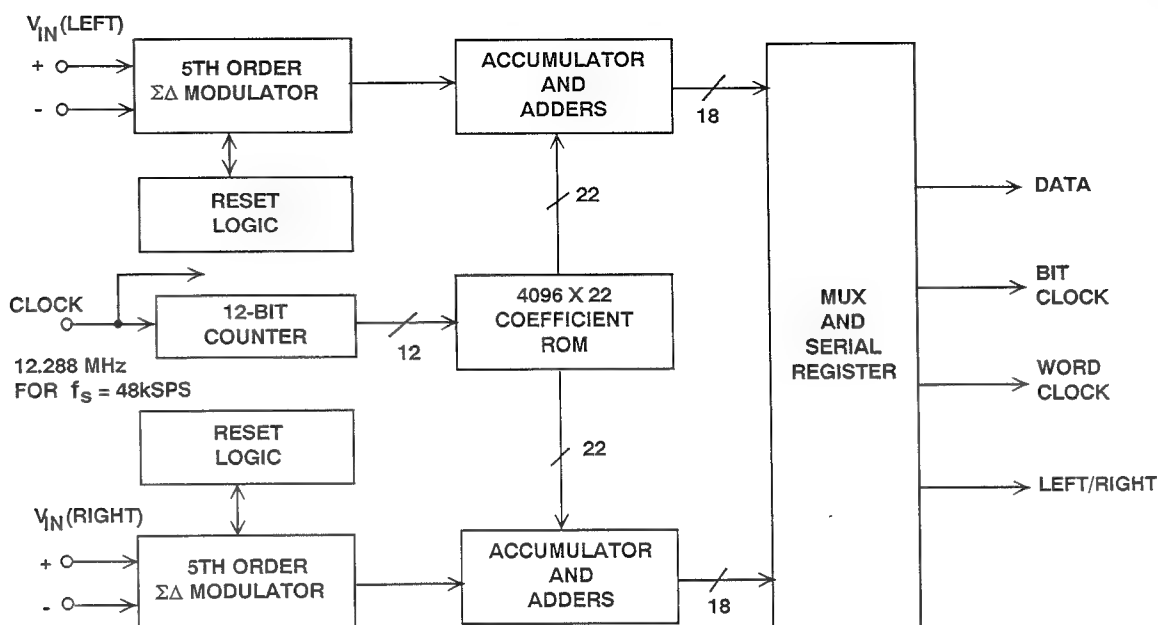


Figure 6.18

AD1879 18 BIT SIGMA-DELTA ADC KEY SPECIFICATIONS

- Two 18 Bit Channels for Stereo Digital Audio
- Interchannel Crosstalk: -110dB at 1kHz
- SNR: 104dB
- THD: 100dB
- Oversampling Ratio: 64x
- Output Word Rate: 55kHz Maximum
- Linear Phase Digital Filter
- Power: 900mW
- 28 Pin, 600-mil Plastic Package

Figure 6.19

AD1879 MODULATOR OUTPUT SPECTRUM




Figure 6.20

For audio ADCs such as the AD1879, the digital lowpass filter cannot be implemented using standard multiply-accumulate structures and present-day semiconductor technology. For example, we require a filter which operates at a sample rate of 3.072MHz ($64 \times 48\text{kHz}$), is flat to 20kHz and has a stopband attenuation of over 115dB starting at 26.2kHz. If we plug these requirements into a standard FIR equiripple design program, the number of coefficients required is 4096. At an output sample rate of 48kHz, we would require a multiply-accumulate time of 5.1ns.

This is clearly too fast for a standard FIR filter structure to implement because of semiconductor process limitations. For this reason, we must use either a parallel processing approach where more than one multiply-accumulate is being executed at any one time, or a multi-rate approach where the decimation is done in more than one step. For the AD1879, a novel parallel processing approach was chosen as described further in Reference 1. The characteristics of this filter are given in Figure 6.21, and the amplitude response in Figure 6.22.

AD1879 DIGITAL FILTER CHARACTERISTICS

- Stopband Attenuation: 118dB
- Passband Ripple: $\pm 0.0008\text{dB}$
- Cutoff Frequency (48kHz output rate): 21.7kHz
- Stopband Frequency (48kHz output rate): 26.2kHz
- Number of Parallel Accumulators: 64 27-bit accumulators
- Coefficient Wordlength: 22bits
- Number of Taps: 4096

Figure 6.21

The AD1879 ADC is a compound monolithic IC. One chip performs the sigma-delta

modulation function, while the second chip performs the digital filtering.

AD1879 DIGITAL FILTER RESPONSE




Figure 6.22

SIGMA-DELTA ADCs FOR LOW FREQUENCY MEASUREMENT APPLICATIONS

Applications such as industrial process control, weigh scales, temperature and pressure measurement instruments require ADCs which can digitize low frequency signals (usually less than 10Hz) to 16 bit or higher precision. In the past, this need was filled almost exclusively by integrating (or dual-slope) ADCs. Sigma-delta converters offer an attractive alternative. In addition to reduced cost and size, low frequency sigma-delta ADCs offer on-board digital filtering as well as system and self-calibration functions. Sampling rates allow processing of signals of

up to 10Hz bandwidths, and power line frequency rejection of sigma-delta ADCs can be maintained over a much wider range of frequency variation than with traditional integrating ADCs. Power supply rejection of dual slope ADCs depends on the instantaneous line frequency variations since the sampling clock is synchronized to the line.

A functional block diagram of the AD7701 monolithic 16 bit sigma-delta ADC is shown in Figure 6.23 and key specifications in Figure 6.24.

AD7701 BLOCK DIAGRAM

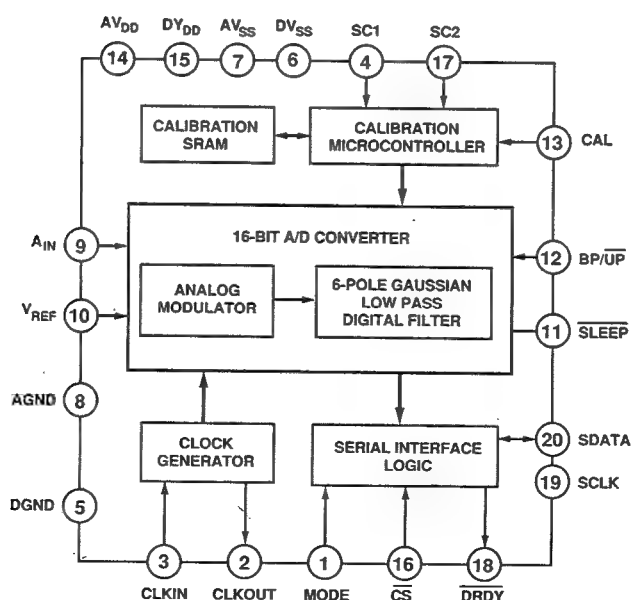


Figure 6.23

AD7701 LOW FREQUENCY MEASUREMENT ADC KEY SPECIFICATIONS

- Monolithic 16 bit ADC
- 0.0015% Linearity Error
- 4 kSPS Output Data Rate
- Programmable Low Pass Filter:
0.1Hz to 10Hz Corner Frequency
- On-Chip Self-Calibration Circuitry
- 0 to +2.5V or ± 2.5 V Input Range
- 40mW Power Dissipation
- 20 μ W Standby Mode
- Flexible Serial Interface
- CS5501 Is a Second Source

Figure 6.24

The AD7701 contains a second-order sigma-delta modulator which samples the analog input signal at a 16kHz rate when the external clock frequency is 4.096MHz. The quantization noise is therefore spread over the bandwidth 0 to 8kHz. The device contains a 6-pole gaussian lowpass digital filter which has a cutoff frequency of 10Hz at the maximum clock rate. The 16kHz sampling rate therefore represents an oversampling ratio of 800 with respect to the 10Hz cutoff frequency. The filter provides 55dB of 60Hz rejection under these conditions. If the clock frequency is halved to give a 5Hz cutoff, 60Hz rejection is better than 90dB. Power supply rejection is 70dB in the 0.1 to 10Hz bandwidth, and PSRR at 60Hz exceeds 120dB due to the digital filter. The frequency response of the digital filter at various clock rates is shown in Figure 6.25.

The long settling time of the internal digital filter (shown in Figure 6.26) in the AD7701 limits its use in multiplexed applications where channels are switched and converted sequentially at high rates. Switching between channels which may have differ-

ent signal levels can cause a step change in the input. The AD7701 is primarily intended for distributed converter systems using one ADC per channel. Multiplexing is possible, provided that sufficient settling time is allowed before data for the new channel is accessed. The gaussian response worst case settling time to $\pm 0.0007\%$ (± 0.5 LSB) is 125ms with a 4.096MHz master clock frequency.

The AD7701 offers two calibration modes using the on-chip calibration microcontroller and SRAM. In the self-calibration mode, zero-scale is calibrated against the analog ground pin (AGND), and fullscale is calibrated against V_{ref} pin. In the system-calibration mode, the AD7701 calibrates its zero and fullscale to voltages present on the analog input pin in two sequential steps, thereby allowing system offsets and/or gain errors to be nulled out.

The AD7703 sigma-delta ADC has a similar architecture to the AD7701, but achieves 20 bits of resolution with 0.0003% linearity error. Key specifications for the AD7703 are summarized in Figure 6.27.

AD7701 DIGITAL FILTER RESPONSE




Figure 6.25

AD7701 DIGITAL FILTER STEP RESPONSE




Figure 6.26

AD7703 LOW FREQUENCY MEASUREMENT ADC KEY SPECIFICATIONS

- Monolithic 22 bit ADC
- 0.0003% Linearity Error
- 4 kSPS Output Data Rate
- Programmable Low Pass Filter:
0.1Hz to 10Hz Corner Frequency
- On-Chip Self-Calibration Circuitry
- 0 to +2.5V or ± 2.5 V Input Range
- 40mW Power Dissipation
- 20 μ W Standby Mode
- Flexible Serial Interface
- CS5503 is a Second Source

Figure 6.27

The AD7710, AD7711, and AD7712 ADCs constitute a family of 21 bit sigma-delta ADCs with on-chip signal conditioning for low frequency, low level measurement applications such as weigh scales, thermocouple temperature measurements, RTD (resistance

temperature detector) temperature measurement, process controllers, and programmable loop controllers. Common features of the three devices are summarized in Figure 6.28, and a block diagram of the AD7710 is shown in Figure 6.29.

AD7710/7711/7712 MEASUREMENT ADC COMMON KEY FEATURES

- 21 Bit Sigma-Delta ADCs, $\pm 0.0015\%$ nonlinearity
- On-Board Differential Input PGA, Gains from 1 to 128
120dB CMR at 50, 60Hz
- First Filter Notch Frequency and Output Data Rate
Programmable from 10Hz to 1kHz
- Sinc³ Filter Response with Cutoff Frequency of 0.262 times
First Filter Notch Frequency
- Ability to Read/Write Calibration Coefficients
- Bidirectional Microcontroller Serial Interface
- Internal/External Reference Option
- Single or Dual Supply Operation
- Low Power (20mW) with Power-Down Mode (10 μ W)

Figure 6.28

AD7710 BLOCK DIAGRAM

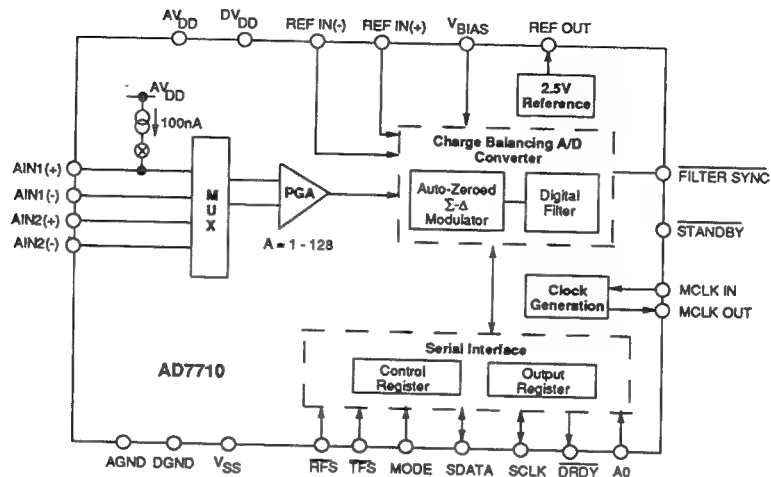


Figure 6.29

AD7710/7711/7712 DIGITAL FILTER RESPONSE




Figure 6.30

An on-board differential input PGA (gain = 1 to 128) enables the user to control full-scale voltage and voltage resolution. The effects of temperature drift are minimized by on-chip self-calibration which removes zero scale and fullscale errors. The internal digital filter has a $(\sin x/x)^3$ response, and 12 bits of data programmed into the control register determine the filter cutoff frequency, the position of the first notch of the filter, and the data rate. In association with the gain selection, it also determines the useful reso-

lution of the device. The first notch frequency (which is also the output data rate) can be programmed from 10Hz to 1kHz. The corresponding -3dB frequency is equal to 0.262 times the first notch frequency. Figure 6.30 shows the filter frequency response for a cutoff frequency of 2.62Hz which corresponds to a first filter notch frequency of 10Hz. The filter response provides greater than 100dB of 50Hz and 60Hz common mode rejection.

Key device-specific features for the three devices are given in Figures 6.31.

AD7710/7711/7712 DEVICE-SPECIFIC FEATURES

AD7710 :

- Two-Channel Differential Low-Level PGA Input

AD7711 :

- Single-Channel Differential Low-Level PGA Input
- RTD (Resistance Temperature Detector) Excitation Current Sources

AD7712 :

- Single-Channel Differential Low-Level PGA Input
- High-Level Analog Input

Figure 6.31

SIGMA-DELTA DACs

Sigma-delta D/A conversion can generally be thought of as the A/D conversion process in the reverse order, where all the basic functions of the digital filter and sigma-delta modulator previously discussed are the same. Sigma-delta DACs offer essentially the same advantages as sigma-delta ADCs. Because of the large oversampling ratio, the requirements on the antialiasing reconstruction filter are greatly relaxed. However, care must be taken to make sure the high frequency noise components contained in the one-bit DAC output are filtered sufficiently.

If a higher order filter is required to reduce this noise, then some of the advantages of the sigma-delta DAC architecture are lost.

Accurate, low-cost, high resolution laser wafer trimmed DACs are readily available, and for this reason there has been less pressure to fully exploit sigma-delta DACs at the component level. The real incentive for developing the sigma-delta DAC technology is because it is the ideal architecture for mixed-signal ICs which require the chip-level integration of the ADC, DAC, and DSP functions.

SIGMA-DELTA DAC CONCEPTS

- Basically a Sigma-Delta ADC in Reverse
- Low-Cost, High Resolution R/2R DACs Proliferate at the Component Level with Oversampling Capability
- Sigma-Delta DACs Ideal for Chip-Level Integration with ADC and DSP Functions
- Antialiasing Filter Must Remove High Frequency Noise

Figure 6.32

The traditional approach to achieving high performance and wide dynamic range using R/2R-based DACs is shown in Figure 6.33. Due to the binary nature of the internal DAC switches, code-dependent transients, or glitches, typically produce some amount of harmonic distortion in the output spectrum. As discussed previously in the DAC section of this seminar, a technique called segmentation can greatly minimize these effects. For the ultimate in spectral purity, the remaining glitches can be removed with a sample-and-hold circuit which holds the DAC output voltage for the duration of the glitch. This technique can eliminate the code-dependent glitches (hence

harmonic distortion) at the expense of introducing some additional energy at the sampling frequency. A lowpass, or smoothing filter is required at the output of the SHA to prevent aliasing as well as eliminate the energy at the sampling rate. The same basic considerations used to define the antialiasing filter used ahead of an ADC apply to the smoothing filter which follows the DAC. For this reason, oversampling relaxes the smoothing filter rolloff requirements in a similar manner. In fact, 2x, 4x, and 8x oversampling techniques are currently in widespread use in compact disk players which use conventional R/2R 16, 18, and 20 bit DACs.

CONVENTIONAL DAC DEGLITCHING

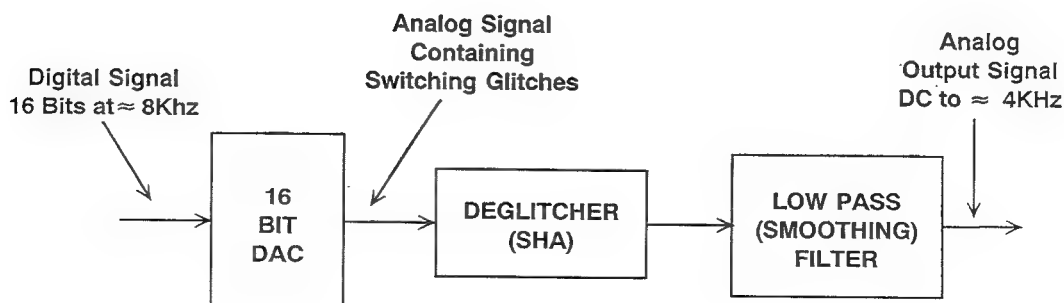


Figure 6.33

The main elements used to implement a sigma-delta DAC are shown in Figure 6.34. The example shown here is for a 16bit DAC which is updated at an 8kHz rate to produce a voiceband output signal having a bandwidth of 4kHz. The 16 bit digital word is fed to a digital interpolation filter where the sampling rate is increased to 1.024MHz, corresponding to an oversampling ratio of 128. This process can be viewed as the reconstruction of a new, higher rate digital signal from an older, lower rate digital signal. Figure 6.35 shows the interpolation of a discrete time signal by a factor of 4. The input signal $x(m)$ is expanded by inserting three zero-valued samples between data samples. The resulting signal $w(m)$ is low-pass filtered to produce $y(m)$ whose sample rate is increased by a factor of 4.

The digital-input sigma-delta modulator noise-shapes the 16-bit 1.024MHz data stream and reduces the sample width to one bit. Unlike the sigma-delta modulator in a sigma-delta ADC, this modulator is all digital. The transfer function is implemented in the digital domain with an IIR filter. This

digital filter performs the same modulator function as in the ADC, where the input signal is effectively lowpass filtered, and the quantization noise is high-pass filtered.

As in the case of a sigma-delta ADC, the 1 bit DAC output is meaningless until it is averaged in some manner. Also, there is a need to remove the shaped quantization noise which resides in the upper frequency area. Finally, there is also a need to reject any images which result about the output Nyquist rate. The analog smoothing filter performs these functions, usually in several stages. It is important for the design of this filter that the filter characteristics match the requirements of the overall system. For example, an audio system would need to have its phase and amplitude response preserved while the output filter also provides the appropriate rejection of higher frequency components. If the smoothing filter is an active filter, care must be taken that the op amps used do not introduce distortion products in the final output due to slewrate limiting and noise.

SIGMA-DELTA DAC

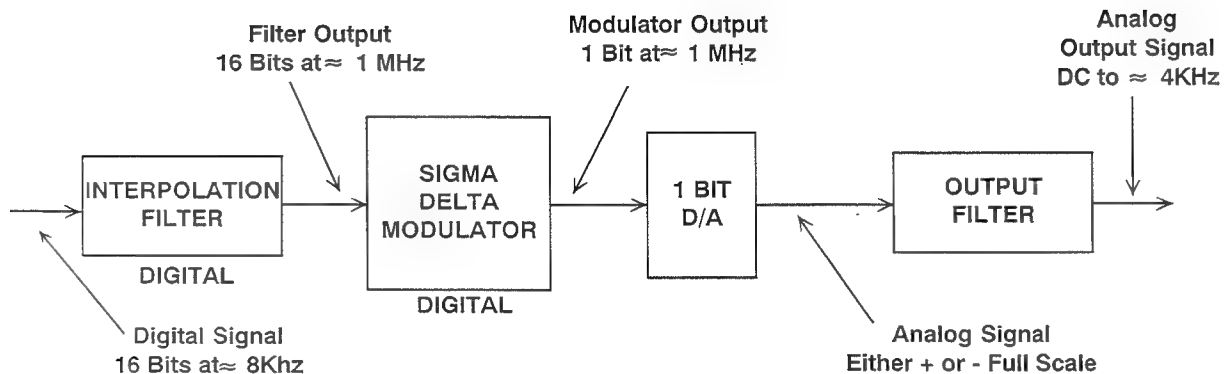


Figure 6.34

INTERPOLATION OF DISCRETE-TIME SIGNAL




Figure 6.35

THE ADSP-28MSP02 SIGMA-DELTA CODEC

The ADSP-28msp02 is a mixed-signal peripheral device available based on sigma-delta design. The device is a linear codec with a 16-bit sigma-delta ADC and DAC, thereby providing a complete analog front

end and back end for high performance voiceband DSP applications. Key features of the IC are summarized in Figure 6.36 and a functional block diagram is shown in Figure 6.37.

KEY FEATURES OF THE ADSP-28msp02 SIGMA-DELTA CODEC

- 16 bit Sigma-Delta ADC
- 16 bit Sigma-Delta DAC
- On-Chip Antialiasing and Smoothing Filters
- 8kSPS Sampling Rate, 128x Oversampling Ratio
- On-Chip Voltage Reference
- 65dB SNR and THD
- Easy Interface to DSP Chips
- 24-pin DIP/SOIC Package
- Single +5V Supply, 100mW Power Dissipation
- Ideal for Voiceband Applications

Figure 6.36

ADSP-28msp02 BLOCK DIAGRAM

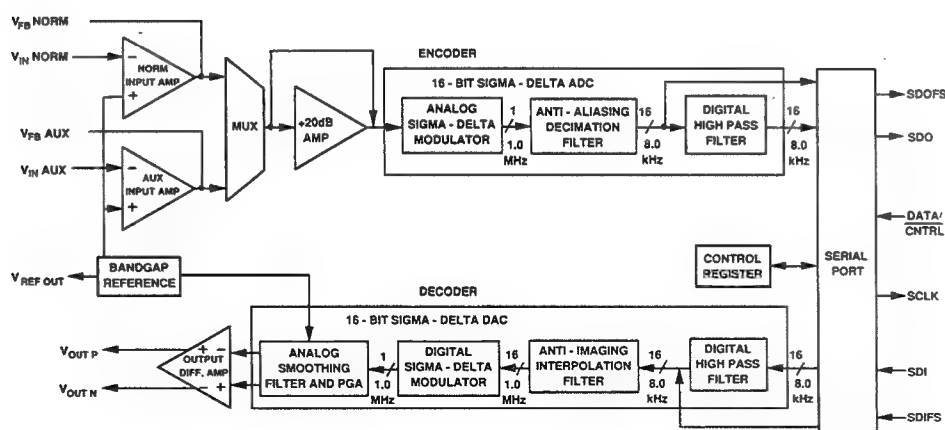


Figure 6.37

6

Compared to traditional m-law and A-law codecs, the ADSP-28msp02's linear coded ADC and DAC maintain wide dynamic range throughout the transfer function. An effective sampling rate of 8kSPS coupled with 65dB SNR and THD performance make the device attractive in many telecommunications applications such as digital cellular telephones. The part is packaged in a 24-pin DIP/SOIC package ensuring a highly integrated and compact solution to voiceband analog processing requirements. The ADSP-28msp02 easily interfaces to the ADSP-2101, ADSP-2105, ADSP-2111, MC56001 and TMS320C25 DSP processors via its serial I/O port; the serial port (SPORT) is used to transmit and receive data or control information to and from the device.

The encoder side of the ADSP-28msp02 consists of two selectable analog input amplifiers and a sigma-delta ADC. The gain of the input amplifiers can be adjusted with the use

of external resistors from -12dB to +26dB. An optional 20dB preamplifier can be inserted before the modulator. The preamplifier and the multiplexer are configured by bits in the control register. The sigma-delta ADC consists of a sigma-delta modulator, an antialiasing decimation filter, and a digital high pass filter. The modulator noise-shapes the signal and produces 1-bit samples at a 1.024MHz rate. This bit stream, representing the analog input, is fed to an antialiasing decimation filter which consists of two low-pass filter stages. The first stage reduces the sampling rate to 40kHz and increases the sample width to 16 bits; the second further reduces the sampling rate to 8kSPS. Each resulting sample is then loaded into the SPORT for transmission.

The decoder consists of a sigma-delta DAC and a differential output amplifier. The DAC reads 16-bit samples at an 8kHz rate from the SPORT. The samples are low- and high-

pass filtered by the digital anti-imaging and high pass filters. The anti-imaging filter interpolates the sampling rate in two stages, first to 40kHz, and then to 1.024MHz. The resulting 16-bit samples are processed by the digital sigma-delta modulator which reduces

the sample width to 1 bit. This bit stream is fed to an analog smoothing filter which converts the data to an analog voltage. The gain of the smoothing filter can be adjusted via the control register from -15dB to +6dB in 3dB steps.

MULTI-STAGE NOISE SHAPING (MASH) SIGMA-DELTA CONVERTERS

As has been discussed, non-linear stabilization techniques have been successfully used to design a fifth-order sigma-delta loop in the AD1879 audio ADC. An alternative approach, called multistage noise shaping (MASH) utilizes cascaded stable first-order loops. Figure 6.38 shows a block diagram of a three-stage MASH ADC. The output of the first integrator is subtracted from the first DAC output to yield the first stage quantization noise, Q_1 . Q_1 is then quantized by the second stage. The output of the second integrator is subtracted from the second DAC output to yield the second stage quantization noise which is in turn quantized by the third stage.

The output of the first stage is summed with a single digital differentiation of the second stage output and a double differentiation of the third stage output to yield the final output. The result is that the quantization noise Q_1 is suppressed by the second stage, and the quantization noise Q_2 is suppressed by the third stage yielding the same suppression as a third-order loop. Since this result is obtained using three first-order loops, stable operation is assured. A comparison of the MASH architecture with the higher-order single-loop architecture is given in Figure 6.39.

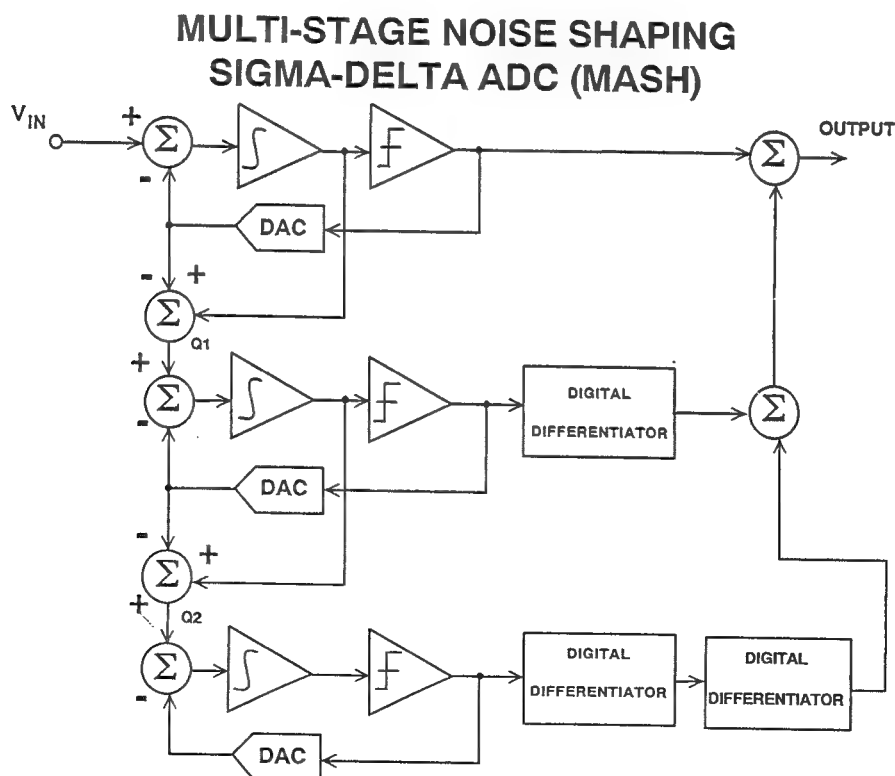


Figure 6.38

MASH TOPOLOGY VERSUS HIGHER-ORDER LOOP SIGMA-DELTA CONVERTERS

- **MASH Cascades Single-Order Loops, therefore Easy to Stabilize**
- **Gain and Phase Matching Critical in MASH Converters for Errors to Cancel**
- **MASH Digital Differentiators Must Match Analog Integrators**
- **Single-Loop Higher Order Modulators Less Subject to Idling Patterns**
- **Single-Loop Higher Order Modulators More Difficult to Understand, Analyze, and Stabilize, But Can Be Done Using Non-Linear Techniques as in AD1879 (5th Order Modulator)**

Figure 6.39

MULTI-BIT SIGMA-DELTA CONVERTERS

So far we have considered only sigma-delta converters which contain a single-bit ADC (comparator) and a single-bit DAC (switch). The block diagram of Figure 6.40 shows a multi-bit sigma-delta ADC which uses an n-bit flash ADC and an n-bit DAC. Obviously, this architecture will give a higher dynamic range for a given oversampling ratio and order of loop filter. Stabilization is easier, since second-order and higher loops can be used. Idling patterns tend to be more

random thereby minimizing tonal effects.

The real disadvantage of this technique is that the linearity depends on the DAC linearity, and thin film laser trimming is required to approach 16-bit performance levels. This makes the multi-bit architecture extremely impractical to implement on mixed-signal ICs. A comparison of the multi-bit versus single-bit sigma-delta converter is given in Figure 6.41.

6

MULTI-BIT FIRST-ORDER SIGMA-DELTA ADC

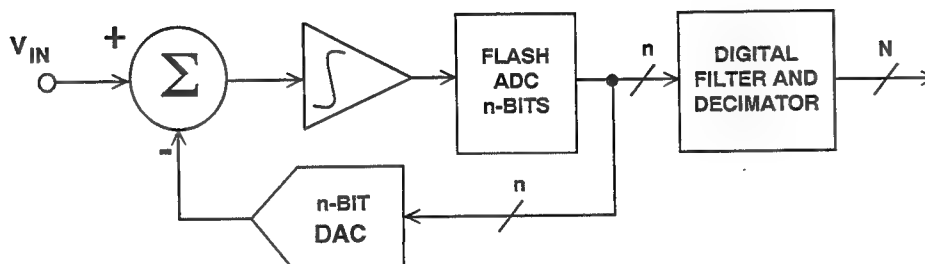


Figure 6.40

MULTI-BIT VERSUS SINGLE-BIT SIGMA-DELTA CONVERTERS

Multi-Bit:

- Higher Dynamic Range for Given Oversampling Ratio and Loop Filter Order
- Higher Order Systems Easier to Stabilize
- Fewer Tonal Effects due to Idling Patterns
- Linearity Depends on DAC
- Thin Film Laser Trimming Required

Single-Bit:

- Perfect Linearity, no Strict Matching Requirements
- No Laser Trimming Required
- Perfect Topology for Mixed-Signal VLSI
- Non-Linear Techniques Required to Stabilize Higher Order Loops (AD1879)

Figure 6.41

SIGMA-DELTA SUMMARY

Although the concepts used in sigma-delta converters are not new by any means, their recent proliferation has been primarily driven by the need for converters which are compatible with mixed-signal VLSI chips. The sigma-delta architecture is ideal for converters for measurement, voiceband, and audio applications. Further exploration of various sigma-delta circuit topologies combined with the development of new processes is sure to push the maximum dynamic range and sampling rates even higher.

It is clear that the sigma-delta converter is not the answer to all data acquisition requirements at the present time. Upper

sampling frequency is limited, thereby excluding video applications, multiplexing inputs is difficult due to the settling time of the internal digital filter, and out-of-range signals may cause saturation of the internal modulators.

On the other hand, the inherently good linearity performance without the need for laser trimming, the relaxation of antialiasing and anti-imaging filter requirements due to oversampling, and the basic sampling nature of the architecture without the need for a SHA will keep sigma-delta development moving at a rapid pace as mixed-signal ICs proliferate.

SIGMA-DELTA SUMMARY

- Inherently Excellent Linearity
- Ideal for Mixed-Signal IC Processes, no Trimming
- No SHA Required
- Upper Sampling Rate Currently Limits Applications to Measurement, Voiceband, and Audio
- Out-of-Range Signals May Cause Modulator Saturation
- Analog Multiplexing Applications Limited by Internal Filter:
Use one Sigma-Delta ADC per Channel!

Figure 6.42

REFERENCES

1. J. Dattorro, A. Charpentier, D. Andreas, The Implementation of a One-Stage Multirate 64:1 FIR Decimator for use in One-Bit Sigma-Delta A/D Applications, **AES 7th International Conference**, May 1989.
2. W.L. Lee and C.G. Sodini, A Topology for Higher-Order Interpolative Coders, **ISCAS PROC. 1987**.
3. P.F. Ferguson, Jr., A. Ganesan and R. W. Adams, One Bit Higher Order Sigma-Delta A/D Converters, **ISCAS PROC. 1990**, Vol. 2, pp. 890-893.
4. R. Koch, B. Heise, F. Eckbauer, E. Engelhardt, J. Fisher, and F. Parzefall, A 12-bit Sigma-Delta Analog-to-Digital Converter with a 15MHz Clock Rate, **IEEE Journal of Solid-State Circuits**, Vol. SC-21, No. 6, December 1986.
5. Wai Laing Lee, A Novel Higher Order Interpolative Modulator Topology for High Resolution Oversampling A/D Converters, **MIT Masters Thesis**, June 1987.
6. D. R. Welland, B. P. Del Signore and E. J. Swanson, A Stereo 16-Bit Delta-Sigma A/D Converter for Digital Audio, **J. Audio Engineering Society**, Vol. 37, No. 6, June 1989, pp. 476-485.
7. R. W. Adams, Design and Implementation of an Audio 18-Bit Analog-to-Digital Converter Using Oversampling Techniques, **J. Audio Engineering Society**, Vol. 34, March 1986, pp. 153-166.
8. B. Boser and Bruce Wooley, The Design of Sigma-Delta Modulation Analog-to-Digital Converters, **IEEE Journal of Solid-State Circuits**, Vol. 23, No. 6, December 1988, pp. 1298-1308.
9. Y. Matsuya, et. al., A 16-Bit Oversampling A/D Conversion Technology Using Triple-Integration Noise Shaping, **IEEE Journal of Solid-State Circuits**, Vol. SC-22, No. 6, December 1987, pp. 921-929.
10. Y. Matsuya, et. al., A 17-Bit Oversampling D/A Conversion Technology Using Multistage Noise Shaping, **IEEE Journal of Solid-State Circuits**, Vol. 24, No. 4, August 1989, pp. 969-975.
11. P. Ferguson, Jr., A. Ganesan, R. Adams, et. al., An 18-Bit 20-kHz Dual Sigma-Delta A/D Converter, **ISSCC Digest of Technical Papers**, February 1991.
12. Steven Harris, The Effects of Sampling Clock Jitter on Nyquist Sampling Analog-to-Digital Converters and on Oversampling Delta Sigma ADCs, **Audio Engineering Society Reprint 2844 (F-4)**, October, 1989.

SECTION VII

DIGITAL SIGNAL PROCESSING TECHNIQUES

DIGITAL SIGNAL PROCESSING TECHNIQUES

■ DIGITAL FILTERING

Finite Impulse Response (FIR) Filters,

The Duality of the Time and Frequency Domains

FIR Filter Implementation in DSP Hardware Using Circular Buffering

FIR Filter Design Techniques

Filter Design Using CAD Techniques

Design Example for an FIR Digital Audio Filter Using CAD Program

Insuring Linear Phase in FIR Filters

Decimation Using FIR Filters

Infinite Impulse Response (IIR) Digital Filters

Summary: FIR Versus IIR Filters

■ FAST FOURIER TRANSFORMS

FFT Hardware Implementation

FFT Design Considerations

Spectral Leakage and Windowing

Data Scaling and Block Floating Point



SECTION VII

DIGITAL SIGNAL PROCESSING TECHNIQUES

DIGITAL FILTERING

Real-time digital filtering is one of the most powerful tools of DSP. Apart from the obvious advantages of virtually eliminating errors in the filter associated with passive component fluctuations over time and temperature, op amp drift (active filters), etc., digital filters are capable of performance specifications that would, at best, be extremely difficult, if not impossible, to achieve with an analog implementation. In addition, the characteristics of a digital filter can be easily changed under software control. Therefore, they are widely used in adaptive-filtering applications such as modems, digital audio, digital mobile radio, and speech processing.

The actual procedure for designing digital filters has the same fundamental elements as that for analog filters. First, the desired filter responses are characterized and the filter parameters are then calculated. Characteristics such as transfer function and phase response are used in the same way. The key difference between analog and digital filters is that instead of calculating resistor, capacitor, and inductor values for an analog filter, coefficient values are calculated for a digital filter. So for the digital filter, numbers replace the physical resistor and capacitor components of the analog filter. These numbers reside in a memory as filter coefficients and are used along with data values from the ADC in performing the filtering calculations.

The digital filter, because it is a discrete function, works with digitized data as opposed to a continuous waveform, and a data point is acquired each sampling period. Because of this discrete nature, we can

reference data samples by numbers such as sample 1, sample 2, sample 3, etc. Figure 7.1, illustrating the basic filtering function, shows a low frequency signal containing higher frequency noise which must be filtered out. This waveform must be digitized with an ADC to produce samples $x(n)$. These data values are fed to the digital filter, which in this case is a lowpass filter. The output data samples, $y(n)$, are used to reconstruct an analog waveform using a DAC.

Digital filters, however, are not the answer to all signal processing filtering requirements. In order to maintain real-time operation, the DSP processor must be able to execute all the steps in the filter routine within one sampling clock period, $1/f_s$. This currently limits their use to primarily voice and audio bandwidth applications. However, it is possible to sacrifice software control and flexibility, and design special hardware digital filters which will operate at video-speed sampling rates. In other cases, the speed limitations can be overcome by first storing the high speed ADC data in a buffer memory. The buffer memory is then read at a rate which is compatible with the speed of the DSP-based digital filter. In this manner, pseudo real-time operation can be maintained as in a radar system, where signal processing is typically done on bursts of data collected after each transmitted pulse. Even in highly oversampled sampled data systems, a simple analog antialiasing filter is usually required ahead of the ADC and after the DAC. Finally, as signal frequencies increase sufficiently, they surpass the capabilities of available ADCs, and digital filtering then becomes impossible, since we no longer have

a sampled data system because we have no ADC. Active analog filtering is not even possible at extremely high frequencies because of op amp bandwidth and distortion limitations, and filtering requirements must

then be met using purely passive components. The primary focus of the following discussions will be on filters which can run in realtime under DSP program control.

DIGITAL FILTERING




Figure 7.1

DIGITAL FILTERING ADVANTAGES

- High Accuracy
- High Performance
- Linear Phase, Constant Group Delay (FIR Filters)
- No Drift Due to Component Variations
- Flexibility, Adaptive Filtering Possible
- Easy to Simulate and Design

Figure 7.2

DIGITAL FILTER LIMITATIONS

- Computation Must be Completed in Sampling Period
- Limited to Voice and Audio Bandwidth Signals if Real-Time Operation is to be Maintained
- Hardwired Digital Filters Required for Video Frequencies
- Analog Filters Still Needed: Antialiasing and High Frequencies
- Lack of High Speed ADCs for Sampling

Figure 7.3

FINITE IMPULSE RESPONSE (FIR) DIGITAL FILTERS

The simplest form of a digital filter is the finite impulse response filter (FIR), and the most elementary form of an FIR filter is a *moving average* filter as shown in Figure 7.4, where we show a 7-day moving average of a dieter's weight plotted along with the daily weights. After 7 days worth of data samples are obtained, the first point on the moving average is computed by adding the 7 data samples together and dividing by 7. Another way to view the process is to *weight* each data sample by a factor of $1/7$ and perform a summation. To obtain the second point on the moving average, the first weighted data sample is subtracted from the summation, and the 8th weighted data sample is added to the summation. This process continues, and can be viewed as a very crude lowpass filtering of the daily readings. The digital implementation of the process is shown in Figure 7.5 which shows the various multiplications, delays, and the summation. The Finite Impulse Response (FIR) filter gets its name because the impulse response is of finite duration; i.e., after seven zero-valued input samples, the filter output goes to zero. When processing an actual electrical signal, a moving average might look like Figure 7.6. It is useful from a mathematical standpoint to view the moving average filter as a *convolution* of the filter impulse response $h(t)$

with the sampled data points $x(t)$ to obtain the output $y(t)$ as shown in Figure 7.7. For a linear convolution, the operation involves multiplying $x(t)$ by a reversed and linearly shifted version of $h(t)$, and then summing the values in the product.

The $\sin(x)/x$ frequency response of the moving average filter is shown in Figure 7.8 for various numbers of taps, N . (Note: in this section N refers to the number of sample points and not the number of bits of resolution of an ADC or DAC!). Note that increasing the number of taps sharpens the rolloff characteristic of the moving average filter but does nothing to improve the undesirable sidelobes.

It is possible to dramatically improve the performance of the simple FIR moving average filter by properly selecting the individual weights or coefficients rather than giving them equal weight. The sharpness of the rolloff can be improved by adding more stages (taps), and the stopband attenuation characteristics can be improved by properly selecting the filter coefficients. The essence of FIR filter design is the appropriate selection of the filter coefficients and the number of taps to realize the desired transfer function $H(f)$. Various algorithms are available to translate the frequency response $H(f)$ into a set of FIR coefficients. Most of this software

SIMPLE MOVING AVERAGE FIR FILTER




Figure 7.4

DIGITAL FORM OF FIR FILTER




Figure 7.5

MOVING AVERAGE FIR FILTER APPLIED TO ANALOG SIGNAL




Figure 7.6

7

MOVING AVERAGE COEFFICIENTS CONVOLVED WITH SAMPLED WAVEFORM




Figure 7.7

FREQUENCY RESPONSE OF MOVING AVERAGE FILTER FOR VARIOUS NUMBER OF TAPS




Figure 7.8

is commercially available and can be run on PCs. The key theorem of FIR filter design is that the coefficients $h(n)$ of the FIR filter are simply the quantized values of the impulse

response of the frequency transfer function $H(f)$. Conversely, the impulse response is the Fourier Transform of $H(f)$.

FACTORS DETERMINING FIR FILTER TRANSFER FUNCTION $H(f)$

- Number of Taps
- Proper Selection of Weighted Filter Coefficients

Figure 7.9

THE DUALITY OF THE TIME AND FREQUENCY DOMAINS

It is useful to digress for a moment and examine the relationship between the time domain and the frequency domain to better understand the principles behind digital filters such as the FIR filter. In a sampled data system, a convolution operation can be

carried out by performing a series of multiplications and accumulations. The convolution operation in the time or frequency domain is equivalent to point by point multiplication in the opposite domain. For example, convolution in the time domain is equivalent to

multiplication in the frequency domain. This is shown graphically in Figure 7.10. It can be seen that filtering in the frequency domain can be accomplished by multiplying all frequency components in the passband by a 1 and all frequencies in the stopband by 0. Conversely, convolution in the frequency domain is equivalent to point by point multiplication in the time domain.

The transfer function in the frequency domain (either a 1 or a 0) can be translated to the time domain by the Fourier transform. This transformation produces an impulse response in the time domain. Since the multiplication in the frequency domain (signal spectrum times the transfer function) is equivalent to convolution in the time domain (signal convolved with impulse response), the signal can be filtered by con-

volving it with the impulse response. The FIR filter is exactly this process. Since it is a sampled data system, the signal and the impulse response are quantized in time and amplitude yielding discrete samples. The discrete samples comprising the impulse response are the FIR filter coefficients.

The mathematics involved in filter design (analog or digital) most always make use of transforms. In continuous-time systems, the Laplace transform can be considered to be a generalization of the Fourier Transform. In a similar manner, it is possible to generalize the Fourier transform for discrete-time sampled data systems, resulting in what is commonly referred to as the z-transform. Details describing the use of the z-transform in digital filter design are given in References 1, 2, and 3.

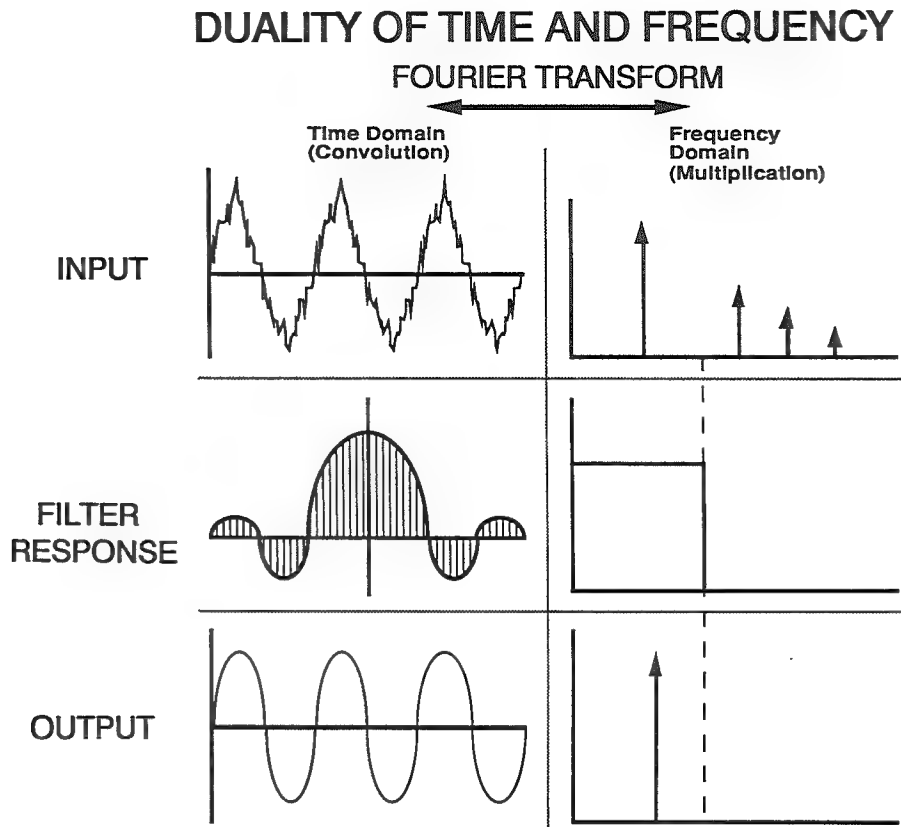


Figure 7.10

FIR FILTER IMPLEMENTATION IN DSP HARDWARE USING CIRCULAR BUFFERING

As has been discussed, an FIR filter (shown in Figure 7.11 must perform the following convolution equation:

$$y(n) = h(n) * x(n) = \sum_{i=0}^{N-1} h(i)x(n-i) \quad ,$$

where $h(i)$ is the filter coefficient array and $x(n-i)$ is the input data array to the filter. The number N , in the equation, represents the number of taps of the filter and relates to the filter performance as has been discussed above.

In the series of FIR filter equations, the N coefficient locations are always accessed sequentially from $h(0)$ to $h(N-1)$. The associated data points circulate through the memory; new samples are added replacing the oldest each time a filter output is computed. A fixed boundary RAM can be used to achieve this circulating buffer effect as shown in Figure 7.12 for a 4 tap FIR filter. The oldest data sample is replaced by the newest after each convolution. A "time history" of the four most recent data samples is kept in RAM.

This delay line can be implemented in fixed boundary RAM in a DSP chip if new data values are written into memory, overwriting the oldest value. To facilitate mem-

ory addressing, old data values are read from memory starting with the value one location after the value that was just written. For example, $x(4)$ is written into memory location 0, and data values are then read from locations 1,2,3, and 0. This example can be expanded to accommodate any number of taps. By addressing data memory locations in this manner, the address generator need only supply sequential addresses regardless of whether the operation is a memory read or write. This data memory buffer is called *circular* because when the last location is reached, the memory pointer must be reset to the beginning of the buffer.

The coefficients are fetched simultaneously with the data. Due to the addressing scheme chosen, the oldest data sample is fetched first. Therefore, the last coefficient must be fetched first. The coefficients can be stored backwards in memory: $h(N-1)$ is the first location, and $h(0)$ is the last, with the address generator providing incremental addresses. Alternatively, coefficients can be stored in a normal manner with the accessing of coefficients starting at the end of the buffer, and the address generator being decremented. In the example shown in Figure 7.12, the coefficients are stored in a reverse manner.

FIR FILTER DESIGN TECHNIQUES

FIR filter design calls for specifying a finite set of N coefficients, $h(n)$, to approximate an idealized filter form. *The filter*

coefficients, $h(n)$, in the time domain correspond to the impulse response of the filter transfer function $H(f)$.

DIRECT FORM FIR FILTER

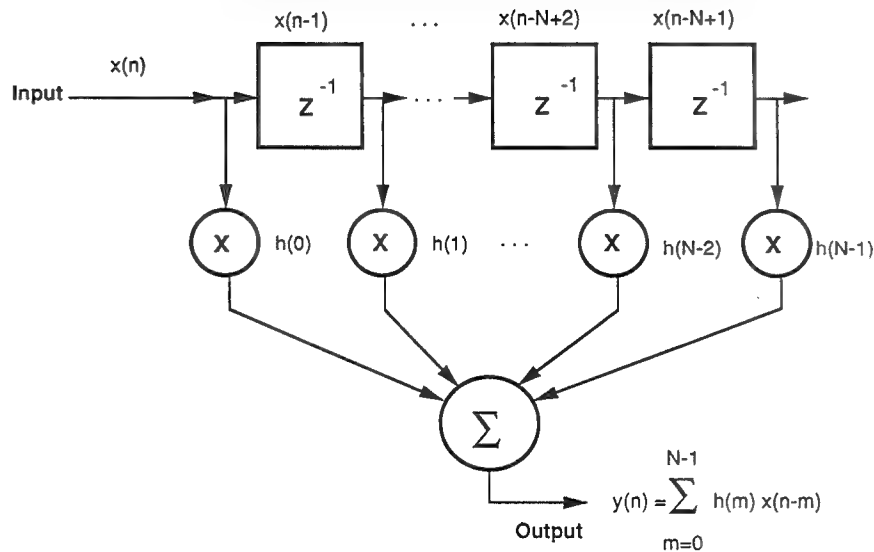
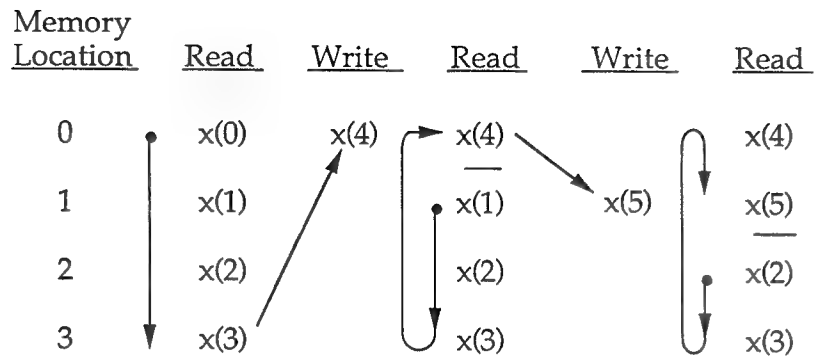


Figure 7.11

DATA MEMORY ADDRESSING FOR 4 TAP FIR FILTER

$$y(n) = h(n) * x(n) = \sum_{i=0}^{N-1} h(i) x(n-i)$$



$$y(3) = h(0)x(3) + h(1)x(2) + h(2)x(1) + h(3)x(0)$$

$$y(4) = h(0)x(4) + h(1)x(3) + h(2)x(2) + h(3)x(1)$$

$$y(5) = h(0)x(5) + h(1)x(4) + h(2)x(3) + h(3)x(2)$$

Figure 7.12

KEY FIR FILTER DESIGN THEOREM

- The Coefficients $h(n)$ of an FIR Filter are Simply the Quantized Values of the Impulse Response of the Frequency Transfer Function $H(f)$
- The Impulse Response is Calculated by Taking the Fourier Transform of $H(f)$

Figure 7.13

In Figure 7.14, 2nd, 4th, and 6th-order ideal Chebyshev lowpass filter transfer functions, optimized for 1dB in-band ripple, are compared with a 91-tap (i.e., 91 coefficients and 91 sequential circular buffer memory locations) digital FIR filter optimized for 0.002dB passband ripple. There is no practical analog equivalent; this is higher order than is realistic with analog hardware (greater than 70 poles using rule-of-thumb approximation). Since the response is flatter within the passband, the signal is reproduced more faithfully, and phase distortion in the

passband is negligible, since all frequencies are delayed equally by the filter. This is another important characteristic of FIR filters (linear phase response and constant group delay) which makes them extremely attractive to digital audio applications.

If the 91-tap FIR filter shown in Figure 7.14 is implemented in the ADSP-2101 microcomputer, each tap requires one processor cycle (80ns). The total processing time is therefore 7.3 μ s. This implies that sampling rates of up to about 136kHz can be achieved and still maintain real-time operation.

91 TAP FIR FILTER RESPONSE COMPARED TO CHEBYSHEV ANALOG FILTER RESPONSE




Figure 7.14

91 TAP FIR FILTER PERFORMANCE CHARACTERISTICS

- 0.002dB Passband Ripple
- Linear Phase
- 80dB Stopband Attenuation
- 136kHz Sampling Rate Possible with ADSP-2101 Processor (80ns Cycle Time per Filter Tap)
- No Analog Equivalent! (70 poles Required!)

Figure 7.15

FIR FILTER DESIGN USING CAD TECHNIQUES

In actual practice, the concepts presented in the above discussions have been implemented in easy to use CAD programs which can be run on most PCs. It is only necessary to specify the desired FIR filter characteristics (sampling frequency, passband frequency,

stopband frequency, passband ripple, and stopband attenuation) as shown in Figure 7.16. The CAD program calculates the number of filter taps required (N), the impulse response, and the filter coefficients.

KEY FILTER DESIGN PARAMETERS




Figure 7.16

FIR FILTER DESIGN CAD PROGRAM INPUTS

- Passband
- Passband Ripple
- Stopband
- Stopband Attenuation
- Wordlength, i.e., 16 Bit Fixed-Point

Figure 7.17

A plot of the frequency response, $H(f)$, along with the impulse response and the step-function response is also available as an output. If the response characteristics are satisfactory, the filter coefficients can then be

downloaded into the DSP processor. The CAD program can also simulate the effects of finite word-length (i.e., performing calculations in 16 bit fixed point arithmetic) on the transfer function.

FIR FILTER DESIGN CAD PROGRAM OUTPUTS

- Frequency Response Plot Showing Effects of Finite Wordlength Arithmetic
- Impulse Response Plot
- Step Function Response
- Number of Taps Required
- Filter Coefficients

Figure 7.18

Other algorithms have been developed for CAD filter designs which optimize the filter performance for various characteristics. An example is the Parks and McClellan program (see Reference 1) which minimizes the maxi-

mum errors between the desired characteristic and the actual characteristic by using the Remez exchange algorithm from approximation theory.

DESIGN EXAMPLE FOR AN FIR DIGITAL AUDIO FILTER USING CAD PROGRAM

For this example, we will design an audio lowpass filter that is designed to operate at a sampling rate of 44.1kHz (standard for CD players). The program is available from Momentum Data Systems, Incorporated (Reference 5). The program is menu-driven and IBM PC compatible. The filter will be imple-

mented as a Direct Form FIR as shown in Figure 7.19.

First, we select the type of filter to be designed from among the Main Menu shown in Figure 7.20. We choose the Equiripple FIR Design (Parks-McClellan)

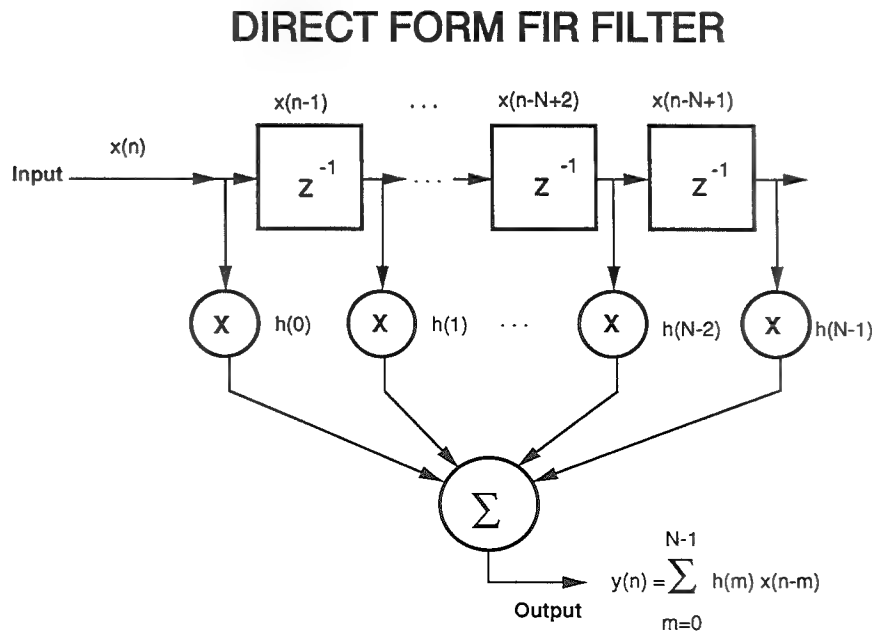


Figure 7.19

FILTER DESIGN AND ANALYSIS SYSTEM MAIN MENU (Screen 1)

- IIR Filter Design
- FIR Filter Design With Windows
- *Equiripple FIR Design (Parks-McClellan)*
- Read Filter Specification File
- System Analysis (Z Domain Input)
- System Analysis (s Domain Input)
- Read System Analysis Input File
- Set System Defaults
- Exit to DOS

Figure 7.20

FINITE IMPULSE RESPONSE FILTER DESIGN MENU (Screen 2)

Filter Type:	1 - Lowpass	4 - Bandstop
	2 - Highpass	5 - Differentiator
	3 - Bandpass	6 - Multiband
Frequency Mode:	H - Hertz	
	R - Radians/Second	
Gain Specification Mode:	1 - Maximum Gain 1.0	
	2 - Nominal Gain 1.0	
Filter Compensation:	Enter X to Select	

Figure 7.21

The second screen then appears as shown in Figure 7.21. This screen is used to select the type of FIR filter (lowpass, highpass, bandpass, etc.) as well as specify the mode for frequency, gain, and whether $\sin(x)/x$ compensation is to be used.

The next screen appears as shown in Figure 7.22 where we enter the sampling

rate, the band edges, and specifications for the passband ripple and stopband attenuation. The example we have chosen is a lowpass filter with a cutoff frequency of 18kHz.

FIR FILTER DESIGN LOWPASS FILTER (Screen 3)

Sampling Frequency:	44100.0
Passband Frequency:	18000.0
Stopband Frequency:	21000.0
Passband Ripple:	1.00000E-02
Stopband Ripple:	96dB
(Attenuation)	

Figure 7.22

The program will then calculate the required filter coefficients. When this calculation is complete, the screen shown in Figure 7.23 appears which lets us know the number of coefficients (taps) required to

implement the filter. If the number of taps is compatible with the throughput of the DSP processor and the sampling rate, the user allows the program to proceed.

FIR DESIGN EXAMPLE (Screen 4)**Estimated Number of Taps of FIR Filter: 69****Enter Number of Taps Desired: 69****Figure 7.23**

If the 69-tap FIR filter is implemented in the ADSP-2101 microcomputer, each tap requires one processor cycle (80ns). The total processing time is therefore 5.5 μ s. This

implies that sampling rates up to about 182kHz can be achieved and still maintain real-time operation.

ADSP-2101 PROCESSOR TIME FOR 69 TAP FIR FILTER

- 80ns (One Processor Cycle) per Tap
- 69 Taps
- 5.5 μ s Processor Time (80ns x 69)
- 182kHz Sampling Rate for Real-Time Operation

Figure 7.24

The next step shown in Figure 7.25 is to quantize the coefficients to the correct number of bits so that the coefficients are compatible with the DSP processor being used. In

this example, the ADSP-2101 is to be used. It is a 16 bit fixed point machine, so the coefficients are quantized to 16 bits.

FIR DESIGN EXAMPLE (Screen 5)**Select the Desired Number of Bits for Quantization****Number of Bits (8 to 32): 16****Figure 7.25**

Now that the coefficients are calculated and properly quantized, we must see what effects on filter performance have been introduced by the quantization process. It should be noted that the filter design program initially calculates the coefficients with very high resolution. When these very accurate coefficients are quantized to a lower resolution, i.e. 16 bits, some accuracy is lost.

This loss in accuracy may adversely affect the performance of the filter. To verify the proper performance, the filter is simulated. In this example the simulation is performed with 16 bit math. Figure 7.26 shows the simulated filter response so that the filter performance can be analyzed. Also available as outputs are the impulse response (shown in Figure 7.27) and the step response (shown

in Figure 7.28). It should be clear that this filter has no analog counterpart. The rule of thumb for calculating the required number of poles of an analog filter having this transition band characteristic (85dB from 18 to 21kHz) would indicate a 65th order filter! (Refer to Section III).

If the filter performance is satisfactory, the coefficient file can be downloaded to the DSP hardware for the filter implementation. If the response is not satisfactory, the design process may be iterated with changes made either to the number of taps or other parameters until the desired response is achieved.

FIR FILTER DESIGN EXAMPLE FREQUENCY RESPONSE




Figure 7.26

FIR FILTER DESIGN EXAMPLE IMPULSE RESPONSE




Figure 7.27

FIR FILTER DESIGN EXAMPLE STEP RESPONSE




Figure 7.28

INSURING LINEAR PHASE IN FIR FILTERS

An advantage of FIR filters is they can always be made to have linear phase response which is a characteristic that makes them extremely attractive in audio and sonar applications. Linear phase means that all input frequencies are delayed by the same amount through the filter. In an FIR filter, this is the time required for the signal to propagate through the N taps. This delay is often referred to as *group delay* when applied

to a band of frequencies. The group delay is constant for a linear phase FIR filter.

In order to insure phase linearity in an FIR filter, it is required that the filter coefficients are symmetric as in the case of a simple lowpass filter (Figure 7.29) or as in the case of a simple highpass filter (Figure 7.30). In addition, using an odd number of taps is also a requirement for linear phase.


**SYMMETRICAL FILTER COEFFICIENTS
PRODUCE LINEAR PHASE RESPONSE -
LOWPASS FILTER**

Figure 7.29


**SYMMETRICAL FILTER COEFFICIENTS
PRODUCE LINEAR PHASE RESPONSE -
HIGHPASS FILTER**

Figure 7.30

DECIMATION USING FIR FILTERS

FIR filters lend themselves to applications where data rate decimation is required, such as in oversampled sigma-delta ADCs. If we want to decimate the output data rate of an FIR filter by a factor of 2, for instance, we would take only every other sample point out of the filter. This also implies that the filter

output computations need only be done every other sampling clock period. In other words, the DSP processor now has two sampling clock intervals to complete the convolution calculation. This implies that either more filter taps can be used, or perhaps a slower processor.

FIR FILTER PROPERTIES SUMMARY

- Always Stable
- Have Linear Phase, Constant Group Delay
- Can be Adaptive
- Low Round-Off Noise
- Computational Advantages When Decimating Output
- Easy to Understand and Implement

Figure 7.31

INFINITE IMPULSE RESPONSE (IIR) DIGITAL FILTERS

As was mentioned previously, digital FIR filters have no real analog counterparts, the closest analogy being the weighted moving average. In addition, FIR filters have only zeros and no poles. On the other hand, IIR filters have traditional analog counterparts (Butterworth, Chebyshev, and Elliptic) and can be analyzed and synthesized using more familiar traditional filter design techniques.

Figure 7.32 shows a second-order lowpass active filter, and its IIR digital filter equivalent is shown in Figure 7.33. This second-order IIR filter is referred to as the *biquad* (because it is described with a biquadratic equation in the z -domain) and forms the basic building block for most higher order IIR designs. The difference equation which describes the characteristics of the filter with 5 coefficients is also shown in the figure.

The general digital filter equation is shown in Figure 7.34 which gives rise to the general transfer function $H(z)$ which contains polynomials in both the numerator and the denominator. The roots of the denominator determine the pole locations of the filter, and

the roots of the numerator determine the zero locations. Although it is possible to construct a high order IIR filter directly from this equation (called the *direct form* implementation), accumulation errors due to quantization errors (finite wordlength arithmetic) may give rise to instability and large errors. For this reason, it is common to cascade several biquad sections with appropriate coefficients rather than use the direct form implementation. The biquads can be scaled separately and then cascaded in order to minimize the coefficient quantization and the recursive accumulation errors. Cascaded biquads execute more slowly than their direct form counterparts, but are more stable and minimize the effects of errors due to finite arithmetic errors. In calculating the throughput time of a particular DSP IIR filter, one should examine the benchmark performance specification for a biquad filter section. For the ADSP-2101, the execution time for a single biquad section is 560ns, corresponding to seven instruction cycles.

SECOND-ORDER ANALOG FILTER IMPLEMENTATION





Figure 7.32

IIR FILTER TOPOLOGY DIRECT FORM I SECOND-ORDER SECTION



$$y(n) = b_0 x(n) + b_1 x(n-1) + b_2 x(n-2) - a_1 y(n-1) - a_2 y(n-2)$$

Figure 7.33

GENERAL FILTER EQUATION

$$y(n) = \overset{\text{FEEDFORWARD}}{\sum_{k=0}^M b_k x(n-k)} + \overset{\text{FEEDBACK}}{\sum_{k=1}^N a_k y(n-k)}$$

GIVES RISE TO THE TRANSFER FUNCTION

$$H(z) = \frac{\sum_{k=0}^M b_k z^{-k} \quad (\text{ZEROS})}{1 - \sum_{k=1}^N a_k z^{-k} \quad (\text{POLES})}$$

Figure 7.34

IIR FILTER PROPERTIES SUMMARY

- Feedback (Recursion)
- Potentially Unstable
- Usually Implemented as Cascaded Biquads Rather than Direct Form
- Non-Linear Phase
- More Efficient Than FIR Filters
- No Computational Advantage when Decimating Output
- Analogous to Analog Filters

Figure 7.35

THROUGHPUT CONSIDERATION FOR IIR FILTERS

- Determine How Many Biquad Sections are Required to Realize the Desired Filter Function
- Multiply by the Execution Time per Biquad (560ns for the ADSP-2101)
- The Result is the Minimum Sampling Period ($1/f_s$) Allowable for Real-Time Operation

Figure 7.36

SUMMARY: FIR VERSUS IIR FILTERS

Choosing between FIR and IIR filter designs can be somewhat of a challenge, but a few basic guidelines can be given. Typically, IIR filters are more efficient than FIR filters because they require less memory and fewer multiplications are needed. IIR filters can be designed based upon previous experience with analog filter designs. IIR filters may exhibit instability problems, but this is much less likely to occur if higher order filters are designed by cascading second-order systems.

On the other hand, FIR filters require more taps and computations for a given

cutoff frequency response, but do exhibit linear phase characteristics. Since FIR filters operate on a finite history of data, if some data is corrupted (ADC sparkle codes, for example) the FIR filter will ring for only $N-1$ samples. Because of the feedback, however, an IIR filter will ring for a considerably longer period of time.

If sharp cutoff filters are needed and processing time is at a premium, IIR elliptic filters are in order. If the number of multiplies is not prohibitive, and linear phase is a requirement, then the FIR should be chosen.

IIR VERSUS FIR FILTERS

IIR FILTERS	FIR FILTERS
More Efficient	Less Efficient
Analog Equivalent	No Analog Equivalent
May be Unstable	Always Stable
Non-Linear Phase Response	Linear Phase Response
More Ringing on Glitches	Less Ringing on Glitches
CAD Design Packages Available	CAD Design Packages Available
No Efficiency Gained by Decimation	Decimation Increases Efficiency

Figure 7.37**FAST FOURIER TRANSFORMS**

In many applications it is desired to process or analyze a signal in the frequency domain. In the analog world, this is easily accomplished using an analog spectrum analyzer. Mathematically, this process can be duplicated by taking the Fourier transform of the continuous-time analog signal. The Fourier transform yields the spectral content of the analog signal. In sampled data systems, however, this process must be accomplished by DSP processing of the ADC output data. Furthermore, there are two

distinct differences between an analog and a digital spectral analysis. First, the output of the ADC is discrete quantized samples of the continuous input, $x(t)$. In sampled data systems, the Discrete Fourier Transform (DFT) performs the transformation of the time domain samples into the frequency domain. In addition, the DFT must operate on a finite number of sampled data points, while the Fourier transform operates on a continuous waveform.

CONTINUOUS AND DISCRETE TIME-TO-FREQUENCY TRANSFORMATIONS

- **Fourier Transform Operates on Continuous-Time Waveforms**
- **Discrete Fourier Transform Operates on a Finite Number of Discrete Time Samples of a Waveform**

Figure 7.38

If $x(n)$ is the sequence of N input data samples, then the DFT produces a sequence of N samples $X(k)$ spaced equally in frequency. The DFT consists of a series of

multiplications and additions where a data word is multiplied by a sinusoid value, and a number of these products are added together as shown in Figure 7.39.

THE DISCRETE FOURIER TRANSFORM (DFT) EQUATION

$$X(k) = \sum_{n=0}^{N-1} x(n)e^{-j2\pi nk/N}, \text{ where}$$

$$e^{-j2\pi nk/N} = \cos(2\pi nk/N) - j\sin(2\pi nk/N)$$

Figure 7.39

7

The DFT can be viewed as a correlation or comparison of the input signal to many sinusoids, evaluating the frequency content of the input signal. For example, a 1024 point DFT would require 1024 samples of the input signal and 1024 points from a sinusoid. Sinusoids of 1024 different frequencies equally spaced from $-f_s/2$ to $+f_s/2$ are used. Each pass of the DFT checks the sinusoid against the input signal to see how much of that frequency is present in the input signal. This is repeated for each of the 1024 frequencies. The result is shown in Figure 7.40 where $N/2$ discrete frequency components appear in the output spectrum. If the sampling frequency is f_s , then the spacing between the spectral lines is f_s/N , or $1/Nt_s$, where t_s is the sampling period, $1/f_s$.

Spectral analysis is most often performed

with complex signals (having both real and imaginary components) so that phase information as well as amplitude and frequency information is obtained. In the above example, 1024 complex data values are multiply/accumulated with 1024 complex sinusoid values. This requires 1024 complex multiplies. This process is repeated for each of the 1024 frequencies for a total of 1024^2 multiplies, or in general terms, N^2 complex multiplies. Even for a powerful DSP device, this number of computations can be cumbersome and time consuming. This amount of computation is only required when all output frequencies are to be calculated. If the value of frequency content for only one or a few frequencies is to be determined, the computational load is not as heavy.

TYPICAL FFT OUTPUTS FOR DIFFERENT RECORD LENGTHS




Figure 7.40

In most spectral analysis situations, however, the entire frequency spectrum up to $f_s/2$ must be computed, so we must find a faster method! The FFT is simply an algorithm to speed up the DFT calculation by reducing the number of multiplications and accumulations required. It was popularized by J. W. Cooley in the 1960s and was actually a rediscovery of an idea of Runge (1903) and Danielson and Lanczos (1942), first occurring prior to the availability of computers and calculators-when numerical calculation could take many manhours.

The FFT is based on taking advantage of certain algebraic and trigonometric symmetries in the DFT computational process. For example, if a 1024 point DFT is performed, 1024^2 (1,048,567) complex multiplications are required. It is possible to break up the 1024 point DFT into two 512 point DFTs and end up with the same results. This is called *decimation*. Each 512 point DFT requires 512^2 (262,144) complex multiplications for a total of 524,288 complex multiplications. This is a significant reduction compared to the original 1,048,567. Figure 7.41 shows an N-point DFT broken up into two N/2-point DFTs. The presence of a phase factor W (sometimes called a *twiddle factor*) on a horizontal line indicates a multiplication by W. The points where the arrows intersect the horizontal lines indicates a summation. The presence of a -1 on the line indicates a sign reversal.

If it's possible to break up the 1024 point DFT into two 512 point DFTs and still get the same result, why can't each 512 point DFT be broken up into two 256 point DFTs for an even greater reduction in computations? Well, they can. This decimation process can continue until the original DFT is broken up into 2 point DFTs (the smallest DFT possible).

The final series of computations, after the decimation process is complete, is the FFT. This is shown for the 8 point DFT in Figure 7.42. Since the FFT was first decimated by a factor of 2, the FFT is known as a Radix-2 FFT. If the *initial* DFT was decimated by a factor of 4, it would be referred to as a Radix-4 FFT. Note that the input data points are taken in normal order, but the outputs are in bit-reversed order. Bit-reversing hardware is therefore common in DSP processors such as the ADSP-2101. The basic calculation, essentially the 2 point DFT, is commonly referred to as a *butterfly* calculation. The FFT is made up of many butterfly calculations. Figure 7.43 shows the basic butterfly for the Radix-2 decimation-in-time FFT which requires one complex multiply operation per butterfly.

The significance of the FFT on the reduction in computations required to do the DFT is shown in Figure 7.44.

FIRST DECIMATION IN TIME OF 8-POINT DFT

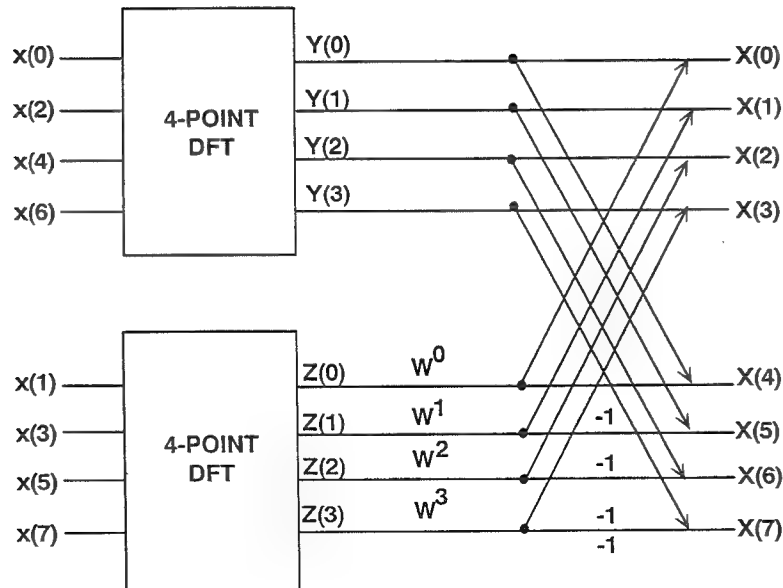


Figure 7.41

7

EIGHT-POINT DECIMATION-IN-TIME FFT NORMAL-ORDER INPUTS, BIT-REVERSED OUTPUTS




Figure 7.42

RADIX-2 DECIMATION-IN-TIME FFT BUTTERFLY




Figure 7.43

COMPUTATIONAL EFFICIENCY OF AN N-POINT FFT

DFT	FFT
N^2 Multiplications	$(N/2)\log_2(N)$
For $N = 1024$	For $N = 1024$
1,048,576 Multiplications	5,120 Multiplications
	200:1

Figure 7.44

Note that the FFT results in the computation of *all* $N/2$ spectral outputs (all or nothing!). If only a few spectral points need to be

calculated, the DFT is more efficient. Calculating a single spectral output using the DFT requires only N complex multiplications.

FFT HARDWARE IMPLEMENTATION

In general terms, the memory requirements for an N point FFT are N locations for real data, N locations for imaginary data, and N locations for the sinusoid data (sometimes referred to as the FFT coefficients or twiddle factors). As long as the memory requirements are met, the DSP processor must perform the necessary calculations in the required time. Many DSP vendors will either give a performance benchmark for a specified FFT size or a calculation time for a butterfly. When comparing FFT specifications, it is important to make sure that the same type of FFT is used in all cases. For

example, a 1024 point FFT benchmark could have been derived from a Radix-2 or Radix-4 FFT and would not be compatible benchmarks since the number of computations required is different.

Once the basic hardware requirements are met, it is the job of the software to make the system realizable. With the same hardware, different software routines make possible a Radix-2, Radix-4, decimation-in-time or decimation-in-frequency algorithm just by manipulating the data in a different manner. An optimized Radix-4 FFT algorithm is given in Reference 6.

DSP FFT HARDWARE BENCHMARK COMPARISONS

- Radix-2, Radix-4 FFT?
- Butterfly Execution Time?
- Total FFT Execution Time?

Figure 7.45

FFT DESIGN CONSIDERATIONS

The first step in designing an FFT is to determine the number of points required, N , or the *record* length. There are several ways to approach this problem. The sampling rate, f_s , must be at least twice the maximum input signal frequency of interest. Once the sampling rate is known, the spectral resolution of the FFT is then given by f_s/N . The more points in the FFT, the better the spectral resolution. This is a prime consideration in spectral analysis applications.

In real-time analysis of speech, for example, the signal bandwidth is approxi-

mately 4kHz, implying a sampling rate of 8kHz. The spectrum of speech is not stationary. The signal must be divided up into windows, T_w , short enough to ensure that individual features are not averaged out in the FFT; all meaning is lost in the *long-term* FFT of speech, for example. But T_w must be long enough to give adequate spectral resolution. It has been determined that for human speech phonemes, 20ms is adequate, hence $T_w = 20\text{ms}$.

REAL-TIME SPEECH ANALYSIS FFT EXAMPLE

- BW = 4kHz, Sampling Rate = 8kHz
- Window = 20ms
- $N > 8\text{kHz} \times 20\text{ms} = 160$, Therefore use $N = 256$
- Can Processor Keep Up?
- ADSP-2101 Benchmark for $N=256$ is 0.59ms
- Yes! With 19.41ms for Other Computations

Figure 7.46

Now, what determines if the FFT can keep up? The number of sample points in the window T_w is equal to $T_w f_s$, or $20\text{ms} \times 8\text{kHz} = 160$ points. This will be rounded up to the nearest power of 2, or 256 points. This says

that the DSP processor must complete the 256 point FFT in less than the data acquisition time per window, T_w . Otherwise real-time processing is not possible, and the computation would have to be done off line. The

ADSP-2101 completes a 256-point FFT in 0.59ms leaving 19.41ms for other computations.

Benchmark FFT processing times for most DSP processors are given by the manufacturer. Figure 7.48 shows Radix-4 benchmark times for the ADSP-2101. The 512-point benchmark time is for a Radix-2 FFT. In evaluating various DSP processors, make sure to compare them under the same condi-

tions. For instance, a Radix-4 FFT is somewhat faster than a Radix-2 FFT.

Figure 7.47 also shows the maximum sampling rates for real-time operation associated with the FFT execution times. These sampling rates indicate that modern DSP microcomputers such as the ADSP-2101 are capable of real-time FFT analysis of signals having bandwidths as great as 100 to 200kHz.

ADSP-2101 BENCHMARK FFT PERFORMANCE AND ASSOCIATED SAMPLING RATES FOR REALTIME OPERATION

FFT SIZE	EXECUTION TIME	MAXIMUM SAMPLING RATE
256	0.59ms	434kHz
512	1.3ms	394kHz
1024	2.9ms	353kHz
2048	6.5ms	315kHz
4096	14.2ms	288kHz

Figure 7.47

FFT OF SINEWAVE HAVING INTEGRAL NUMBER OF CYCLES IN WINDOW




Figure 7.48

SPECTRAL LEAKAGE AND WINDOWING

Spectral leakage in FFT processing can best be understood by considering the case of performing an FFT on a pure sinewave input. Two conditions will be considered. In Figure 7.48, the ratio between the sampling rate and the input sinewave frequency is such that precisely an integral number of cycles are contained within the data window (or record length). This results in a single tone FFT spectral response at the sinewave frequency as shown in the figure. Figure 7.49 shows the condition when the sinewave does not contain an integral number of cycles within the data window. The discontinuities at the endpoints are equivalent to multiplying the sinewave by a rectangular windowing pulse which has a $\sin(x)/x$ frequency domain response. The discontinuities in the time domain result in leakage in the frequency domain, because many spectral terms are needed to fit the discontinuity. Because of the endpoint discontinuity, the FFT spectral response shows the main lobe of the sinewave being smeared, and a large number of associated sidelobes which have the basic characteristics of the rectangular time pulse.

Since in practical FFT spectral analysis applications the exact frequencies are unknown, something must be done to minimize these sidelobes. This is done by choosing a windowing function other than the rectangular window. The input time samples are multiplied by an appropriate windowing function which brings the signal to zero at the edges of the window. The selection of an appropriate windowing function is primarily a tradeoff between main-lobe spreading and sidelobe rolloff. Leakage can also be reduced by padding the data with zeros and performing a correspondingly longer FFT. Reference 4 is highly recommended for an in-depth look at windows.

The time-domain and frequency-domain characteristics of a simple windowing function (the Hanning Window) are shown in Figure 7.50. A comparison of the frequency response of the Hanning window and the more sophisticated Minimum 4-Term Blackman-Harris window is given in Figure 7.51.

FFT OF SINEWAVE HAVING NON-INTEGRAL NUMBER OF CYCLES IN WINDOW




Figure 7.49

TIME AND FREQUENCY REPRESENTATION OF HANNING WINDOW




Figure 7.50

DATA SCALING AND BLOCK FLOATING POINT

The results of the butterfly calculation can be larger than the inputs to the butterfly. This data growth can pose a potential problem in a DSP with a fixed number of bits. To prevent data overflow, the data needs to be scaled before hand, leaving enough extra bits for growth. Alternatively, the data can be scaled after each pass of the FFT. The technique of scaling data after each pass of the FFT is known as *block floating point*. It

is called this because the full array of data is scaled as a block regardless of whether or not each element in the block needs to be scaled. The complete block is scaled so that the relative relationship of each data word remains the same. For example, if each data word is shifted right one bit (divided by 2), the absolute values have been changed but relative to each other, the data stays the same.

COMPARISON OF WEIGHTING FUNCTIONS

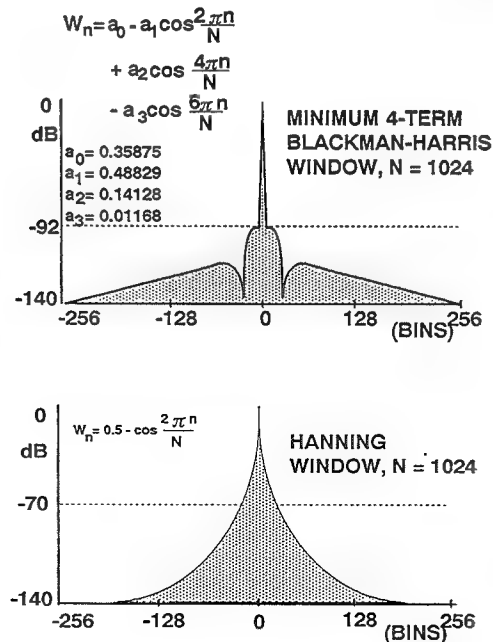


Figure 7.51

FFT SUMMARY

- The FFT is an Algorithm, not an Approximation
- Computational Speed is not Achieved at the Expense of Accuracy
- The FFT is a Fast Implementation of the DFT
- Resolution of the FFT in Frequency is f_s/N , N = Record Length
- Endpoint Discontinuities in Time Usually Require Smoothing Using Windowing Functions
- Real-Time FFT Processing Possible at Sampling Rates in Excess of 100kHz Using DSP Microcomputers

Figure 7.52

REFERENCES

1. Richard J. Higgins, **Digital Signal Processing in VLSI**, Prentice-Hall, 1990.
2. A. V. Oppenheim and R. W. Schaffer, **Digital Signal Processing**, Prentice-Hall, 1975.
3. L. R. Rabiner and B. Gold, **Theory and Application of Digital Signal Processing**, Prentice-Hall, 1975.
4. Fredrick J. Harris, On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform, **Proc. IEEE**, Vol. 66, No. 1, 1978 pp. 51-83.
5. Momentum Data Systems, Inc. Costa Mesa, CA.
6. Fares Eidi, An Optimized Radix-4 Fast Fourier Transform (FFT), **Analog Devices Application Note E1329-5-9/89**. Available from Analog Devices.
7. **High Speed Design Seminar**, Analog Devices, 1990.
8. Amy Mar, Editor, **Digital Signal Processing Applications Using the ADSP-2100 Family**, Prentice-Hall, 1990.
9. C. S. Williams, **Designing Digital Filters**, Prentice-Hall, 1986.
10. R. W. Ramirez, **The FFT: Fundamentals and Concepts**, Prentice-Hall, 1985.

SECTION VIII

DSP HARDWARE

DSP HARDWARE

- **RISC VS. CISC VS. DSP ARCHITECTURES**
- **RISC AND DSP APPLICATIONS**
- **DSP PROCESSOR REQUIREMENTS**
 - Fast Arithmetic
 - Zero Overhead Looping
 - Extended Dynamic Range
 - Dual Operand Fetch
 - Circular Buffers
- **ADSP-2101 MICROCOMPUTER GENERAL DESCRIPTION**
- **ADSP-2101 ARCHITECTURE OVERVIEW**
 - Arithmetic Logic Unit (ALU)
 - Multiplier/Accumulator (MAC)
 - Shifter
 - Data Address Generators (DAGs)
 - Program Sequencer
 - Serial Ports
 - System Interface
- **DEVELOPMENT SYSTEM**

SECTION VIII

DSP HARDWARE

RISC VERSUS CISC ARCHITECTURES

As central processor (CPU) architectures developed, their instruction sets became more sophisticated. The complex-instruction-set (CISC) computer includes instructions for basic processor operations, plus single instructions that are highly sophisticated—for example, to evaluate a high-order polynomial. But CISC has a price: many of the instructions execute via microcode in the CPU and require numerous clock cycles—plus silicon real estate for code storage.

In contrast, the reduced-instruction-set computer (RISC) recognizes that, in many applications, basic instructions such as LOAD and STORE—with simple addressing schemes—are used much more frequently than the advanced instructions, and should not incur an execution penalty. These simpler instructions are hardwired in the CPU logic to execute in a single clock cycle, reducing execution time and CPU complexity.

RISC AND DSP APPLICATIONS

Although the RISC approach offers many advantages in general purpose computing, it is not well suited to DSP. For example, most RISCs do not support single-instruction multiplication, a very common and repetitive operation in DSP. The DSP is optimized to accomplish its task fast enough to maintain real-time operation in the context of the application, which requires single-cycle arithmetic operations and accumulations.

DSP algorithms have unique requirements not found in general purpose computing such as circular buffering, pointer updating and fast looping with zero overhead, bit reversing, barrel shifting, scaling, and data-dependent execution branching. Each of these should execute within the DSP instruction, and not as a separate time-consuming instruction cycle. The computational unit within the DSP must run efficiently, with

data arriving from at least two separate data memories with no time penalty for data access. CISCs and RISCs support virtually none of these needs.

Similarly, software programming for RISCs and CISCs differs from that used in DSP. RISCs and CISCs are programmed in high-level languages to minimize software development and hide the assembly language from the programmer. For real-time DSP applications, however, code optimization (primarily of execution time, but also of memory usage) requires that the software engineer use assembly language to get satisfactory performance. If the initial results are not satisfactory after simulation, critical sections of the program are examined and recoded if necessary to reduce overall execution time.

ARCHITECTURES

- Complex Instruction Set Computer (CISC)
- Reduced Instruction Set Computer (RISC)
- DSP Processor

Figure 8.1

DSP PROCESSOR REQUIREMENTS

Many DSP algorithms (such as digital filters and FFTs) rely heavily on the efficient performance of the straight-forward sum-of-products mathematics. Examining the

equation shown in Figure 8.2 reveals that there are some fundamental properties implicit in the mathematics.

THE DSP KERNEL EQUATION

$$\Sigma = h(0)x(0) + h(1)x(1) + h(2)x(2) + \dots + h(N-1)x(N-1)$$

- The Equation is Based on Multiply-Accumulates (MACs)
- There are N MACs - One for Each product
- Each Product is Formed from 2 Values, One Value is a Signal, $x(i)$, the Other a Stored Coefficient, $h(i)$

Figure 8.2

The three algorithmic properties place specific requirements on processor architectures aimed at digital signal processing. The fundamental properties of the DSP kernel

function, $\Sigma h(i)x(i)$, and effects of finite word size arithmetic (i.e. quantization errors) combine to produce five DSP architectural requirements shown in Figure 8.3.

DSP ARCHITECTURE REQUIREMENTS

- Fast Arithmetic
- Zero Overhead Looping
- Extended Dynamic Range
- Dual Operand Fetch
- Circular Buffers

Figure 8.3

FAST ARITHMETIC

Fast arithmetic is the simplest of these requirements to understand. Since real-time DSP applications are driven by performance, the multiply-accumulate or MAC time is a central requirement; faster MACs mean potentially higher bandwidth. It is critical to remember that MAC time alone does not define DSP performance. This often forgotten fact leads to an inadequate measure of processor performance by simply examining its MIPS (million instructions per second) rating. Since most DSP and DSP-like architectures feature MACs that can execute an

instruction every cycle, most processors are given a MIPS rating equal to its MAC throughput. This does not necessarily account for the other factors that can degrade a processors overall performance in real-world applications. The other four criteria can wipe out MAC gains if they are not satisfied.

In addition to the requirement for fast arithmetic, a DSP should be able to support other general purpose math functions and should therefore have an appropriate arithmetic logic unit (ALU) and a shifter function.

ZERO OVERHEAD LOOPING

Zero overhead looping is required by the repetitive nature of the kernel equation. The multiply-accumulate function and the data fetches needed are repeated N times every time the kernel function is calculated. Traditional microprocessors implement loops that have one instruction execution time or more

of overhead associated with them. DSP architectures provide hardware support that eliminates the need for looping instructions within the loop body. For true DSP architectures, the difference of zero overhead body looping and programmed looping can easily exceed 20% in available bandwidth.

EXTENDED DYNAMIC RANGE

Extended dynamic range is a requirement of finite word size arithmetic. The basic convolution requires repeated addition. Since each product can be a full range value, one accumulation of two products can cause overflow beyond the word size of the processor. Traditional microprocessors and microcomputers address overflow by providing an overflow flag. The program then has the added burden of testing the flag and conditionally adjusting the results. Architectures

designed specifically for DSP reduce this problem by providing extended precision in the accumulator function of the MAC. Simply adding 8 bits extends performance by about 48dB. A further refinement of extended precision accumulation allows signed overflow and underflow so that the intermediate values of the MAC can track the real-world values that are applied with minimal chance of loss of accuracy.

DUAL OPERAND FETCH

Regardless of the nature of a processor, performance limitations are generally based on bus bandwidth. In the case of general purpose processors, code is dominated by single memory fetch instructions usually

addressed as base plus offset value. This leads architects to embed fixed data into the instruction set so that this class of memory access is fast and memory efficient. DSP on the other hand is dominated by instructions

requiring two independent memory fetches. This is driven by the basic form of the convolution $\sum h(i)x(i)$. The goal of fast dual operand fetches is to keep the MAC fully loaded. We saw in the discussion on MACs that the performance of a DSP is first limited by MAC time. Assuming an adequate MAC cycle time, two data values need to be supplied at the same rate; reductions in operand fetch bandwidth will result in corresponding reductions in MAC bandwidth. Ideally, the operand fetches occur simultaneously with

the MAC instruction so that the combination of the MAC and memory addressing occurs in one cycle.

Dual operand fetch is implemented in DSPs by providing separate buses for program memory data and data memory data. In addition, separate program memory address and data memory address buses are also provided. The MAC can therefore receive inputs from each data bus simultaneously. This architecture is often referred to as the Harvard Architecture.

CIRCULAR BUFFERS

If we examine the kernel equation more carefully, the advantages of circular buffering in DSP applications become apparent. A Finite Impulse Response (FIR) filter is used to demonstrate the point. First, coefficients or tap values for FIR filters are periodic in nature. Second, FIR filters use the newest real-world signal value and discard the oldest value.

In the series of FIR filter equations, the N coefficient locations are always accessed sequentially from $h(0)$ to $h(N-1)$. The associated data points circulate through the memory; new samples are added replacing the oldest data each time a filter output is computed. A fixed boundary RAM can be used to achieve this circulating buffer effect. The oldest data sample is replaced by the newest after each convolution. A "time history" of the N most recent samples is kept in RAM.

This delay line can be implemented in fixed boundary RAM in a DSP chip if new data values are written into memory, overwriting the oldest value. To facilitate memory addressing, old data values are read from memory starting with the value one location after the value that was just written. In a 4-tap FIR filter, for example, $x(4)$ is written into memory location 0, and data values are then read from locations 1, 2, 3, and 0. This example can be expanded to accommodate any number of taps. By addressing data memory locations in this manner, the ad-

dress generator need only supply sequential addresses regardless of whether the operation is a memory read or write. This data memory buffer is called *circular* because when the last location is reached, the memory pointer must be reset to the beginning of the buffer.

The coefficients are fetched simultaneously with the data. Due to the addressing scheme chosen, the oldest data sample is fetched first. Therefore, the last coefficient must be fetched first. The coefficients can be stored backwards in memory: $h(N-1)$ is the first location, and $h(0)$ is the last, with the address generator providing incremental addresses. Alternatively, coefficients can be stored in a normal manner with the accessing of coefficients starting at the end of the buffer, and the address generator being decremented.

This allows direct support of unit delay taps without software overhead. These data characteristics are DSP algorithm-specific and must be supported in hardware to achieve the best DSP performance. Implementing circular buffers in hardware allows buffer parameters (i.e. start, length, etc.) to be set up outside of the core instruction loop. This eliminates the need for extra instructions within the loop body. Lack of a hardware implementation for circular buffering can significantly impact MAC performance.

SUMMARY

Any processor can accomplish any software task, given enough time. However, DSPs are optimized for the unique computational requirements of real-time, real-world signal processing. Traditional computers are

better suited for tasks that can be performed in non-real-time. In the following section, we will examine the architecture of a high-performance DSP Microcomputer, the ADSP-2101.

ADSP-2101 MICROCOMPUTER GENERAL DESCRIPTION

The ADSP-2101 is a single-chip microcomputer optimized for digital signal processing and other high-speed numeric processing applications. It combines the complete ADSP-2100 core architecture (three computational units, data address generators, and a program sequencer) with two serial ports, a programmable timer, extensive interrupt capabilities and on-board program and data

memory RAM. The ADSP-2101 has 1K words of (16 bit) data memory RAM and 2K words of (24 bit) program RAM on chip.

The ADSP-2101's flexible architecture and comprehensive instruction set support a high degree of operational parallelism. In one cycle the ADSP-2101 can perform the functions shown in Figure 8.4.

ADSP-2101 SINGLE-INSTRUCTION CYCLE CAPABILITY

- Generate the Next Program Address
- Fetch the Next Instruction
- Perform One or Two Data Moves
- Update One or Two Data Address Pointers
- Perform a Computational Operation
- Receive and Transmit Data Via the Two Serial Ports
- Update Timer

Figure 8.4

ADSP-2101 ARCHITECTURE OVERVIEW

Figure 8.5 is an overall block diagram of the ADSP-2101. The processor contains three independent computational units: the ALU, the multiplier/accumulator (MAC) and the shifter. The computational units process 16-bit data directly and have provisions to support multiprecision computations. The ALU performs a standard set of arithmetic and logic operations; division primitives are also supported. The MAC performs single-cycle multiply, multiply/add, and multiply/

subtract operations. The shifter performs logical and arithmetic shifts, normalization, denormalization, and derive exponent operations. The shifter can be used to efficiently implement numeric format control including multiword floating-point representations.

The internal result (R) bus directly connects the computational units so that the output of any unit may be the input of any unit on the next cycle.

A powerful program sequencer and two dedicated data address generators ensure efficient use of these computational units. The sequencer supports conditional jumps, subroutine calls and returns in a single cycle. With internal loop counters and loop stacks, the ADSP-2101 executes looped code with zero overhead; no explicit jump instructions are required to maintain the loop.

The data address generators (DAGs) handle address pointer updates. Each DAG keeps track of four address pointers. When-

ever the pointer is used to access data (indirect addressing), it is post-modified by the value of a specified modify register. A length value may be associated with each pointer to implement automatic modulo addressing for circular buffers. With two independent DAGs, the processor can generate two addresses simultaneously for dual operand fetches. The circular buffering feature is also used by the serial ports for automatic data transfers.

ADSP-2101 BLOCK DIAGRAM

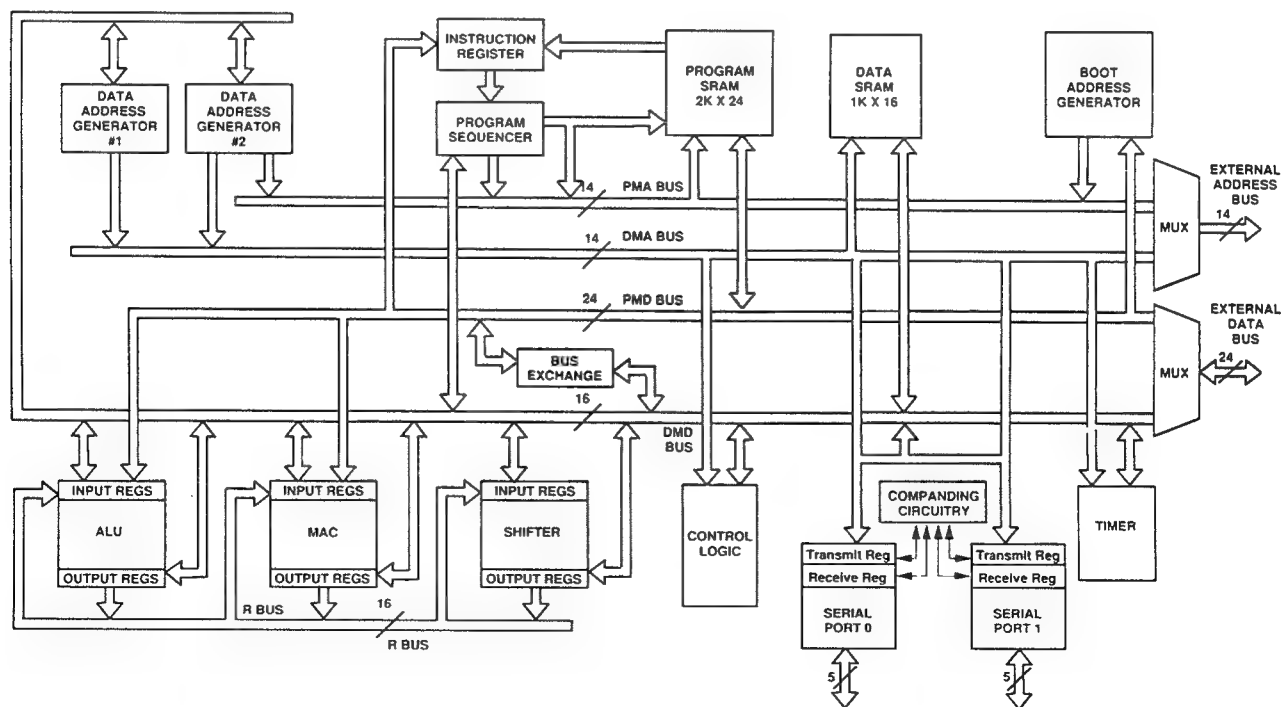


Figure 8.5

COMMON FEATURES OF ADSP-2100 FAMILY

- Arithmetic Logic Unit
- Multiplier/Accumulator (With 40-Bit Accumulator)
- Barrel Shifter
- Two Data Address Generators
- Program Sequencer

Figure 8.6

Efficient data transfer is achieved with the use of five internal buses. The two address buses (PMA and DMA) share a single external address bus, and the two data

buses (PMD and DMD) share a single external data bus. The $\overline{\text{BMS}}$, $\overline{\text{DMS}}$, and $\overline{\text{PMS}}$ signals indicate which memory space the external buses are being used for.

ADSP-2101 INTERNAL BUSES

- Program Memory Address (PMA) Bus
- Program Memory Data (PMD) Bus
- Data Memory Address (DMA) Bus
- Data Memory Data (DMD) Bus
- Result (R) Bus

Figure 8.7

Program memory can store both instructions and data, permitting the ADSP-2101 to fetch two data operands in a single cycle, one from program memory and one from data memory as well as an instruction from program memory. Because the on-board program memory is so fast, the ADSP-2101 can fetch an operand from program memory and the next instruction in the same cycle.

The memory interface supports slow memories and memory-mapped peripherals with programmable wait state generation. External devices can gain control of buses with bus request/grant signals ($\overline{\text{BR}}$ and $\overline{\text{BG}}$). One execution mode allows the ADSP-2101 to continue running while the buses are granted to another master as long as an external memory operation is not required.

The other execution mode requires the processor to halt while the buses are granted.

The two serial ports provide a complete serial interface with companding in hardware and a wide variety of framed and frameless data transmit and receive modes of operation. Each port can generate an internal programmable serial clock or accept an external serial clock.

Boot circuitry provides for loading on-chip program memory automatically from byte-wide external memory. After reset three wait states are automatically generated. This allows, for example, an 80ns ADSP-2101 to use an external 250ns EPROM as boot memory. Multiple programs can be selected and loaded from the EPROM with no additional hardware.

ARITHMETIC LOGIC UNIT (ALU)

The ALU is shown in Figure 8.8. The ALU provides a standard set of arithmetic and logic functions: add, subtract, negate, increment, decrement, absolute value, AND, OR, Exclusive OR and NOT. Two divide primitives are also provided. The ALU takes two 16-bit inputs, X and Y, and generates one 16-bit output, R. The carry-in feature enables multiword computations. Six arithmetic status bits are generated: AZ (zero), AN (negative), AV (overflow), AC (carry), AS (sign), and AQ (quotient).

The X input port can be fed by either the AX register set or any result register via the R-bus (AR, MR0, MR1, MR2, SR0, or SR1). The AX register set contains two registers, AX0 and AX1. The AX registers can be loaded from the DMD bus. The Y input port can be fed by either the AY register set or the ALU feedback (AF) register. The AY register set contains two registers, AY0 and AY1. The AY registers can be loaded from either the DMD bus or the PMD bus.

The register outputs are dual-ported so that one register can provide input to the ALU while either one simultaneously drives the DMD bus. The ALU output can be loaded into either the AR register or the AF register.

The AR register has a saturation capability; it can be automatically set to plus or minus the maximum value if an overflow or underflow occurs. The AR register can drive both the R bus and the DMD bus and can be loaded from the DMD bus.

The ALU contains a duplicate bank of registers shown in Figure 8.8 behind the primary registers. The secondary set contains all the registers described above (AX0, AX1, AY0, AY1, AF, AR). Only one set is accessible at a time. The two sets of registers allow fast context switching, such as for interrupt servicing.

ALU BLOCK DIAGRAM

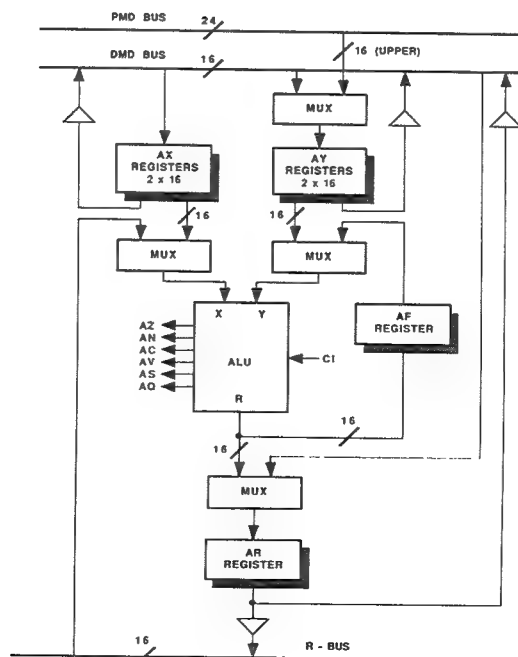


Figure 8.8

ALU FEATURES

- Feedback Paths
- Six Status Flags
- Saturation
- Implements Divide Primitives
- Provisions for Double Precision
- Complete Set of Background Registers

Figure 8.9

EXAMPLE ALU INSTRUCTIONS

- $AR = AX0 + AY0$
- $AF = MR1 \text{ XOR } AY1$
- $AR = AX0 + AF$

Figure 8.10

MULTIPLIER/ACCUMULATOR (MAC)

The multiplier/accumulator (MAC) implements high-speed multiply, multiply/add, and multiply/subtract operations. A block diagram of the MAC section is shown in Figure 8.11.

The multiplier takes two 16-bit inputs, X and Y, and generates one 32-bit output, P. The 32-bit output is routed to a 40-bit accumulator which can add or subtract the P output from the value in MR. MR is a 40-bit register which is divided into three sections: MR0 (Bits 0-15), MR1 (Bits 16-31), and MR2 (Bits 21-29). The result of the accumulator is either loaded into the MR register or into the 16-bit MAC feedback (MF) register. The multiplier accepts the X and Y inputs in either signed or unsigned formats.

In default operation the result is shifted one bit to the left to remove the redundant sign bit for fractional justification; an optional mode on the ADSP-2101 inhibits this shift for integer operations. The accumulator generates one status bit, MV, which is set

when the accumulator result overflows the 32-bit boundary. A saturate instruction is available to change the contents of the MR register to the maximum or minimum 32-bit value if MV is set. The accumulator also has the capability for rounding the 40-bit result at the boundary between bit 15 and bit 16.

The MAC and ALU registers are similar. The X input port can be fed by either the MX register set (MX0, MX1) or any result register via the R-bus (AR, MR0, MR1, MR2, SR0, or SR1). The MX register set is readable and loadable from the DMD bus and has dual-ported outputs.

The Y input port can be fed by either the MY register set (MY0, MY1) or the MF register. The MY register set is readable from the DMD bus and readable and loadable from both the DMD and the PMD bus. Its outputs are also dual-ported. The accumulator output can be loaded into either the MR register or the MF register. The MR register is connected to both the R-bus and

MAC BLOCK DIAGRAM

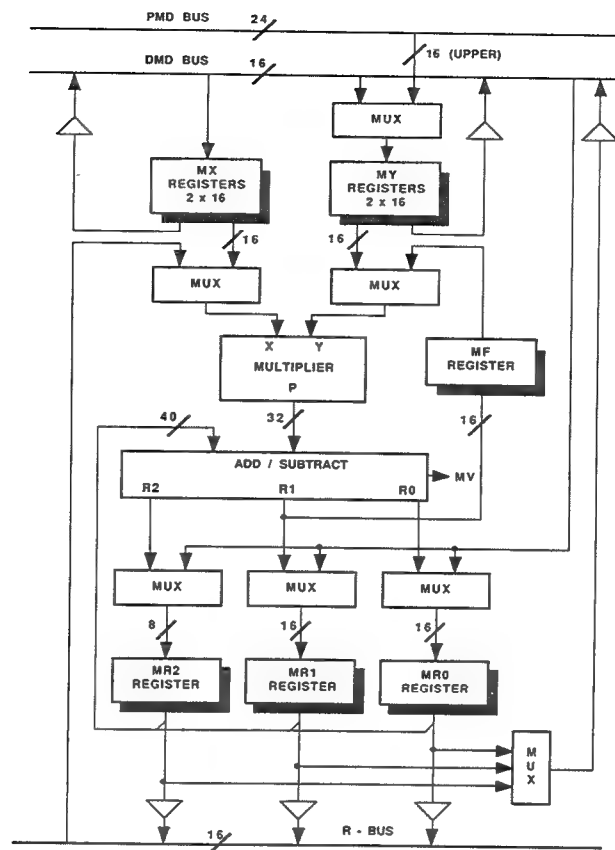


Figure 8.11

the DMD-bus. Like the ALU section, the MAC section contains two complete banks of

registers (MX0, MX1, MY0, MY1, MF, MR0, MR1, MR2) to allow fast context switching.

MAC FEATURES

- Feedback Paths
- 40-Bit Accumulator
- Saturation
- Mixed Mode Input Operands
- Provisions for Multiprecision
- Complete Set of Background Registers

Figure 8.12

EXAMPLE MAC INSTRUCTIONS

- $MR = MX0 * MY0$
- $MR = 0$
- $MF = AR * MF$
- $MR = MX0 * MF$
- $MR = MR + MX1 * MY0$

Figure 8.13

SHIFTER

The shifter gives the ADSP-2101 its unique capability to handle data formatting and numeric scaling. Figure 8.14 shows a block diagram of the shifter.

The shifter can be divided into the following components: the shifter array, the OR/PASS logic, the exponent detector and the exponent compare logic. These components give the shifter its six basic functions: arithmetic shift, logical shift, normalization, denormalization, derive exponent and derive block exponent.

The shifter array is a 16×32 barrel shifter. It accepts a 16-bit input and can place it anywhere in the 32-bit output field, from off-scale right to off-scale left. The shifter can perform arithmetic shifts (shifter output is sign-extended to the left) or logical shifts (shifter output is zero-filled to the left). The placement of the 16-bit input is determined by the control code (C) and the HI/LO reference signal.

SHIFTER BLOCK DIAGRAM

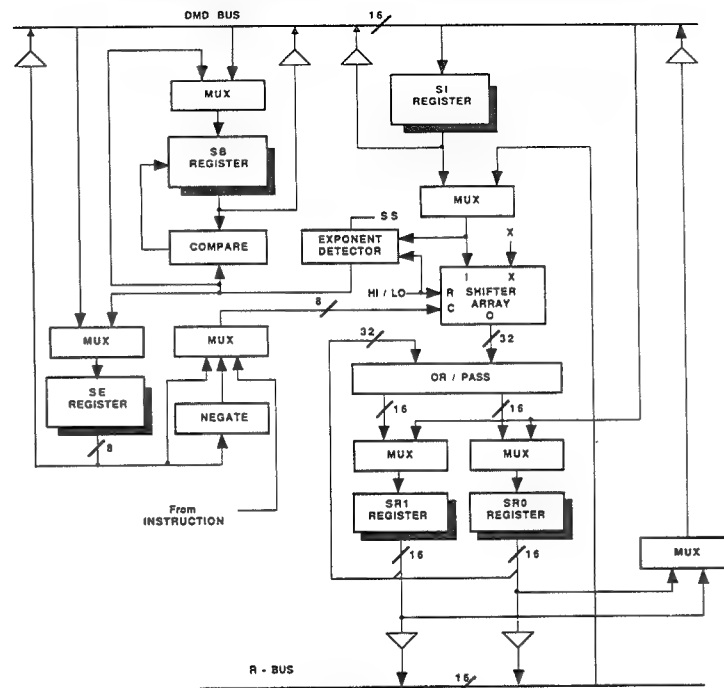


Figure 8.14

SHIFTER FEATURES

- **Arithmetic and Logical Shifts**
- **Left and Right Shifts**
- **True Block Floating Point**
- **Direct Support for Double Precision**
- **Complete Set of Background Registers**

Figure 8.15

SHIFTER OPERATIONS

- **Normalize**
- **Denormalize**
- **Shift Immediate**
- **Derive Exponent**
- **Derive Block Exponent**

Figure 8.16

EXAMPLE SHIFTER INSTRUCTIONS

- **SR = ASHIFT SI BY -6**
- **SR = LSHIFT SR BY 3**
- **SR = NORM MR1**

Figure 8.17

DATA ADDRESS GENERATORS (DAGs)

A block diagram of a data address generator is shown in Figure 8.18. The data address generators (DAGs) provide indirect addressing for data stored in the program and data memory spaces. The processor contains two independent DAGs so that two data operands (one in program memory and one in data memory) can be addressed simultaneously. The two data address generators are identical except that DAG1 has a bit reversal option on the output (used for FFTs)

and can only generate data memory addresses, while DAG2 can generate both program and data memory addresses but has no bit reversal capability. Both DAGs can also be used for serial port autobuffering.

There are three register files in each DAG: the modify (M) register file, the index (I) register file, and the length (L) register file. Each of these register files contains four 14-bit registers which are readable and loadable from the DMD bus. The I registers hold the

DATA ADDRESS GENERATOR BLOCK DIAGRAM

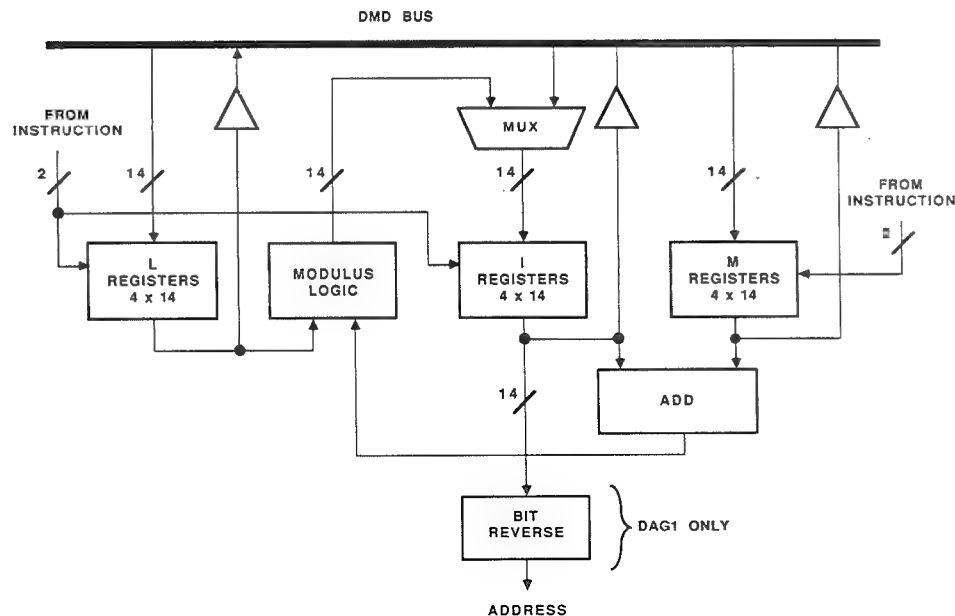


Figure 8.18

actual addresses used to access external memory. When using the indirect addressing mode, the selected I register content is driven onto either the PMA or DMA bus. This value is post-modified by adding the (signed) contents of the selected M register. The modified address is passed through the modulus logic.

Associated with each I register is an L register which contains the length of the

buffer addressed by the I register. The L register and the modulus logic together enable circular buffer addressing with automatic wraparound at the buffer boundary. Automatic wraparound is also used by the serial ports to generate the serial port interrupt when operating in autobuffering mode. The modulus logic is disabled by setting the L register to zero.

ADDRESS GENERATOR FEATURES

- Automatic Modulo Addressing
- Simultaneous Address Update
- Bit-Reverser (DAG # 1)

Figure 8.19

EXAMPLE ADDRESSING INSTRUCTIONS

- **AX0 = DM (I0, M3)**
- **MODIFY (I1, M2)**
- **MR = MR+MX0*MY0, MX0=DM(I0,M1), MY0=PM(I4,M4)**

Figure 8.20**PROGRAM SEQUENCER**

The program sequencer incorporates powerful and flexible mechanisms for program flow control such as zero-overhead looping, single-cycle branching (both conditional and unconditional), and automatic interrupt processing. Figure 8.21 shows a block diagram of the program sequencer. The sequencing logic controls the flow of the program execution. It outputs a program memory address onto the PMA bus from one of four sources: the PC incrementer, PC stack, instruction register, or interrupt controller. The next address source selector controls which of these four sources are selected based on the current instruction word and the processor status. A fifth possible source for the next program memory address is provided by DAG2 when a register indirect jump is executed.

The program counter (PC) is a 14-bit register which contains the address of the currently executing instruction. The PC output goes to the incrementer. The incremented output is selected as the next program memory address if program flow is sequential. The PC value is pushed into the 16 x 14 PC stack when a CALL instruction is

executed or when an interrupt is processed. The PC stack is popped when the return from a subroutine or interrupt is executed. The PC stack is also used in zero-overhead looping.

The program sequencer section contains six status registers. These are the Arithmetic Status register (ASTAT), the Stack Status register (SSTAT), the Mode Status register (MSTAT), the Interrupt Control register (ICNTL), the Interrupt Mask register (IMASK) and the Interrupt Force and Clear register (IFC).

The interrupt controller allows the processor to respond to the six possible interrupts with a minimum of overhead. Individual interrupt requests are logically ANDed with the bits in IMASK; the highest priority unmasked interrupt is then selected.

The interrupt control register, ICNTL, allows each interrupt to be set as either edge- or level-sensitive. Depending on Bit 4 in ICNTL, interrupt routines can either be nested with higher priority interrupts taking precedence or processed sequentially with only one interrupt service active at a time.

ADSP-2101 PROGRAM SEQUENCER

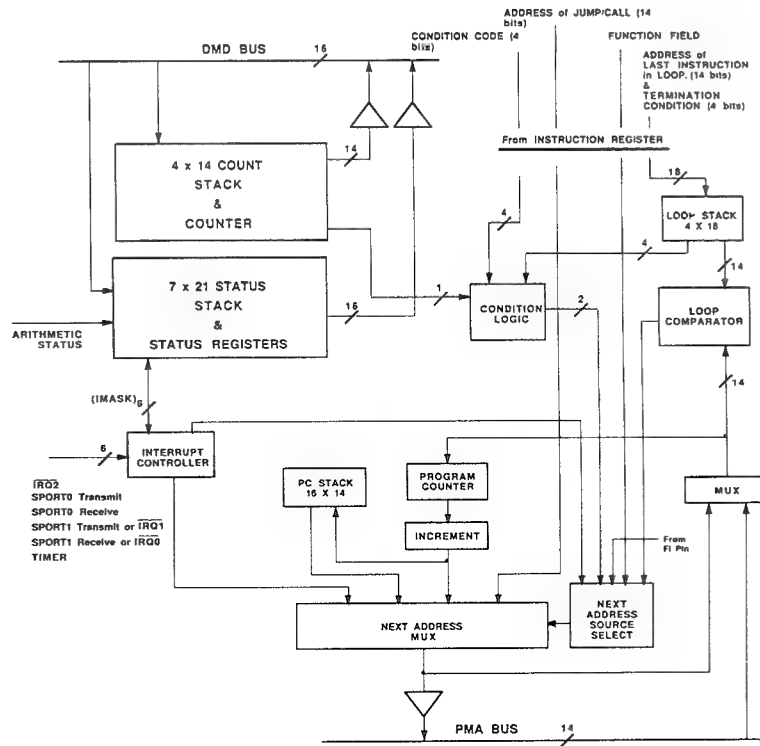


Figure 8.21

PROGRAM SEQUENCER FEATURES

- Automatic Operation, Transparent to User
- Full Interrupt Capabilities
- Four Stacks
- Single-Cycle Conditional Branch
- Zero-Overhead Looping

Figure 8.22

SERIAL PORTS

The ADSP-2101 incorporates two complete serial ports (SPORT0 and SPORT1) for serial communications and multiprocessor coordination. A block diagram of one of the serial ports is shown in Figure 8.23.

Each serial port has a 5-pin interface consisting of the signals shown in Figure 8.24.

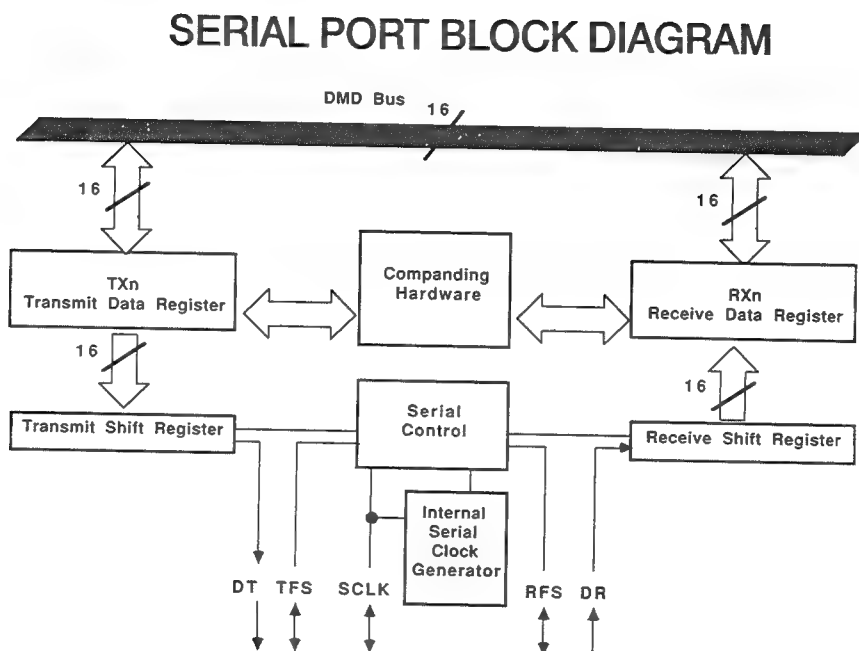


Figure 8.23

SERIAL PORT INTERFACE LINES

■	SCLK	Serial Clock I/O
■	RFS	Receive Frame Synch I/O
■	TFS	Transmit Frame Synch I/O
■	DR	Serial Data Receive
■	DT	Serial Data Transmit

Figure 8.24

Each SPORT has a receive and a transmit register. Companding (a contraction of COMpressing and exPANDING) is the process of logarithmically encoding data to reduce the number of bits that must be sent. The

ADSP-2101 supports both of the widely used algorithms for companding: A-law and μ -law. The type of companding can be independently selected for each SPORT.

SERIAL PORT FEATURES

- Dual Purpose Function of Serial Port 1
- Optional μ -Law and A-Law Companding
- Automatic Data Memory Buffering
- Programmable Word Length
- Multichannel Capabilities

Figure 8.25

SYSTEM INTERFACE

Figure 8.26 shows a basic system configuration with the ADSP-2101, two serial codecs, a boot EPROM and optional external Program and Data memories. Up to 15K words of data memory and 16K words of program memory can be supported. Pro-

grammable wait state generation allows the processor to interface easily to slow memories.

The ADSP-2101 also provides one external interrupt and two serial ports or three external interrupts and one serial port.

ADSP-2101 BASIC SYSTEM CONFIGURATION

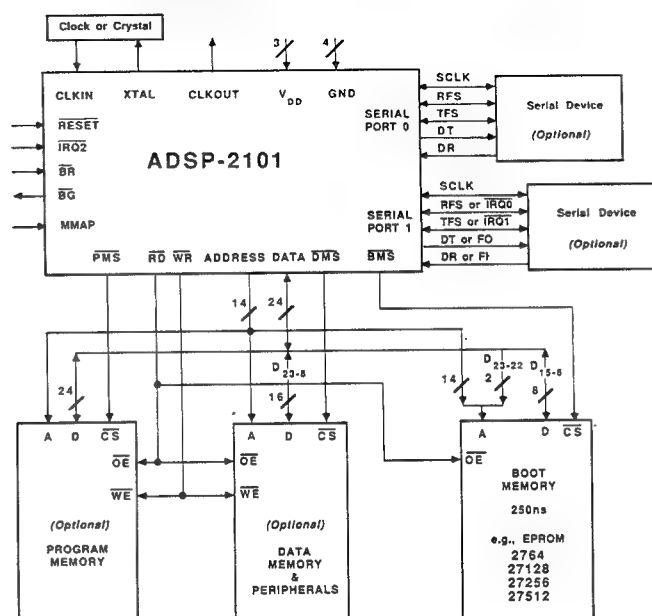


Figure 8.26

DEVELOPMENT SYSTEM

The ADSP-2101 is supported by a complete set of tools for software and hardware system development. The System Builder provides a high-level method for defining the architecture of systems under development. The Assembler produces the object code and the Linker combines object modules and library calls into an executable file. The

Simulator provides an interactive instruction-level simulation with a reconfigurable user interface. A PROM Splitter generates PROM burner compatible files. The C Compiler generates ADSP-2101 assembly source code. An Emulator aids in the hardware debugging of ADSP-2101 systems.

HARDWARE AND SOFTWARE DEVELOPMENT TOOLS

- System Builder
- C Compiler
- Assembler
- Linker
- Simulator
- Prom Splitter
- Evaluation Board
- In-Circuit Emulator

Figure 8.27

ANALOG DEVICES DSP PROCESSOR PORTFOLIO

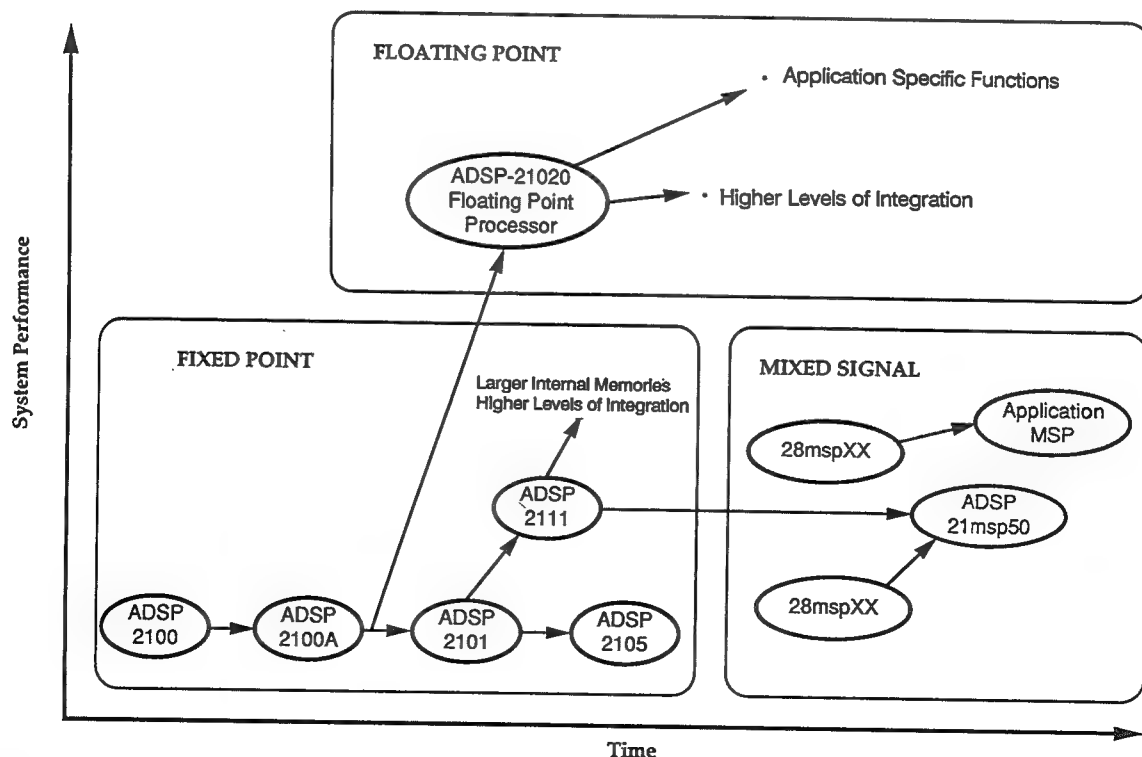


Figure 8.28

REFERENCES

(AVAILABLE FROM ANALOG DEVICES)

1. ADSP-2100/ADSP-2100A Digital Signal Processor, Data Sheet
2. ADSP-2101 DSP Microcomputer, Data Sheet
3. ADSP-2105 DSP Microcomputer, Data Sheet
4. ADSP-2111 DSP Microcomputer with Host Port, Data Sheet
5. ADDS-21XX DSP Software Development Tools, Data Sheet
6. ADDS-21XX DSP Hardware Development Tools, Data Sheet
7. ADDS-2101-SW DSP Software Development Tools, Data Sheet
8. ADSP-2101 Emulator, Data Sheet
9. ADDS-2101-EZ Tools, Data Sheet
10. ADSP-2101 User's Manual
11. ADSP-2101 Cross-Software Manual
12. ADSP-2101 Emulator Manual
13. ADSP-2101 EZ-ICE Manual
14. ADSP-2101 EZ-LAB Manual
15. ADSP-2111 User's Manual
16. *Digital Signal Processing Applications Using the ADSP-2100 Family*
(Applications Handbook, Volumes 1, 2, and 3)
17. ADSP-2100 Family Applications Handbook, Volume 4

SECTION IX

INTERFACING ADCs AND DACs TO DIGITAL SIGNAL PROCESSORS

INTERFACING ADCs AND DACs TO DIGITAL SIGNAL PROCESSORS

- **PARALLEL INTERFACING TO DSP PROCESSORS: READING
DATA FROM MEMORY-MAPPED PERIPHERAL ADCs**

Parallel ADC to DSP Interface

- **PARALLEL INTERFACING TO DSP PROCESSORS: WRITING
DATA TO MEMORY-MAPPED DACs**

Parallel DAC to DSP Interface

- **SERIAL INTERFACING TO DSP PROCESSORS**

Serial ADC to DSP Interface

Serial DAC to DSP Interface

- **INTERFACING I/O PORTS AND CODECS TO DSPs**

- **SERIAL VERSUS PARALLEL DSP INTERFACE SUMMARY**

SECTION IX

INTERFACING ADCs AND DACs TO DIGITAL SIGNAL PROCESSORS

As the technology in the rapidly growing field of Mixed Signal Processing evolves, more highly integrated DSP products are being introduced (such as the ADSP-21msp50) which contain on-chip ADCs and DACs as well as the DSP, thereby eliminating most component-level interface problems. Stand-alone ADCs and DACs are now available with interfaces especially designed for DSP chips, thereby minimizing or eliminat-

ing external interface support or *glue* logic. High performance sigma-delta ADCs and DACs are currently available in the same package (called a codec or coder/decoder) such as the ADSP-28msp02. These products are also designed to require minimum glue logic when interfacing to the most common DSP chips. This section discusses the various data transfer and timing issues associated with the various interfaces.

PARALLEL INTERFACING TO DSP PROCESSORS:

READING DATA FROM MEMORY-MAPPED PERIPHERAL ADCs

Interfacing an ADC or a DAC to a fast DSP parallel bus (such as the ADSP-2101, ADSP-2100, or the TMS320C25) requires an understanding of how the DSP processor reads data from a memory-mapped peripheral (the ADC) and how the DSP processor writes data to a memory-mapped peripheral (the DAC). We will first consider some general timing requirements for reading and writing data.

A block diagram of a typical parallel DSP interface to an external ADC is shown in Figure 9.1. This diagram has been greatly simplified to show only those signals associated with *reading* data from an external memory-mapped peripheral device. The timing diagram for the ADSP-2101 read-cycle is shown in Figure 9.2.

The *read* process begins when the peripheral device (such as an ADC) asserts the *processor interrupt request line* ($\overline{\text{IRQ}}$). The processor then places the address of the peripheral initiating the interrupt request on the *memory address bus* (A0 - A13). At the same time, the processor asserts the *data*

memory select line ($\overline{\text{DMS}}$). The two internal address buses of the ADSP-2101 (program memory address bus and data memory address bus) share a single external address bus, and the two internal data buses (program memory data bus and data memory data bus) share a single external data bus. The *boot memory select* (BMS), *data memory select* (DMS), and *program memory select* (PMS) signals indicate which memory space the external buses are being used for. These signals are typically used to enable an external address decoder as shown in Figure 9.1. The output of the address decoder drives the *chip select* input of the peripheral device.

The *memory read* ($\overline{\text{RD}}$) is asserted t_{ASR} ns after the $\overline{\text{DMS}}$ line is asserted. The sum of the address decode delay plus the peripheral chip select setup time should be less than t_{ASR} in order to take full advantage of the time the $\overline{\text{RD}}$ low-time. The $\overline{\text{RD}}$ line remains low for t_{RP} ns. The *memory read* signal is used to enable the tri-state parallel data outputs of the peripheral device. The $\overline{\text{RD}}$ line is connected to the appropriate pin

ADC/ADSP-2101 PARALLEL INTERFACE

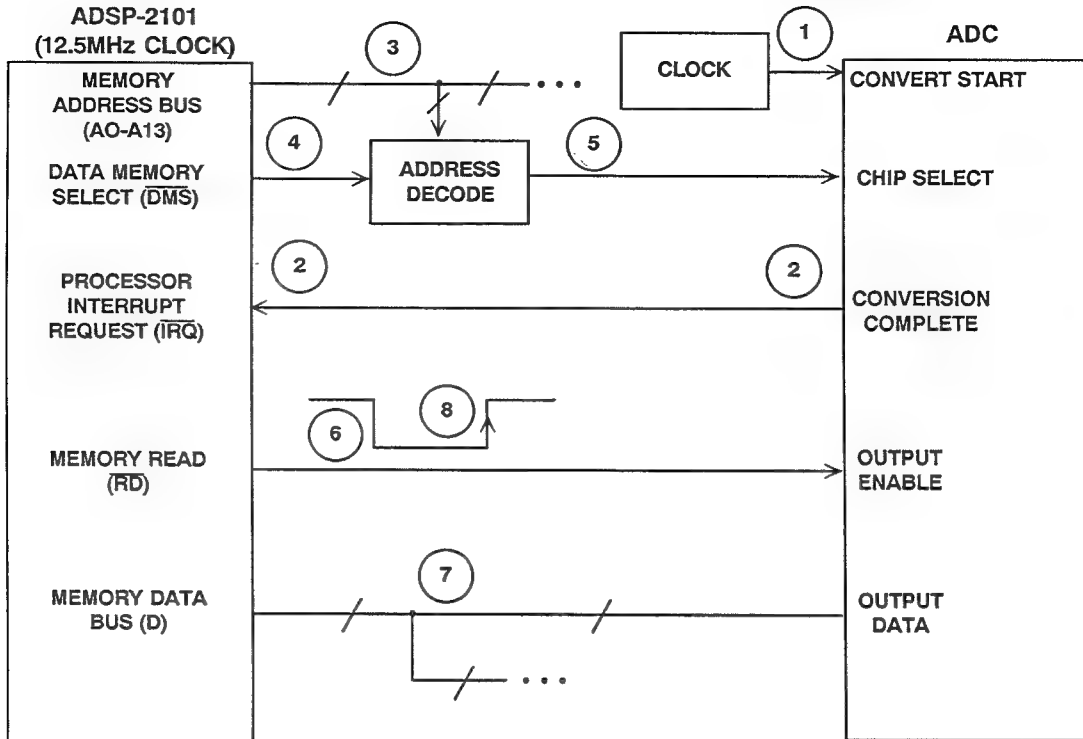


Figure 9.1

ADSP-2101 MEMORY READ TIMING




Figure 9.2

on the peripheral device usually called *output enable* or *read*. The rising edge of the \overline{RD} signal is used to clock the data on the data bus into the DSP processor. After the rising edge of the \overline{RD} signal, the data on the data

bus must remain valid for t_{RDH} ns, the data hold time. In the case of the ADSP-2101, this value is 0ns.

The key timing requirements for the peripheral device are shown in Figure 9.3.

PARALLEL PERIPHERAL DEVICE READ INTERFACE KEY REQUIREMENTS

- Peripheral Device Data Outputs Must Be Tri-State
- Address Decode Delay Plus Peripheral Chip Select Setup Time Must Be Less Than Address and Data Memory Setup Times (5ns min for ADSP-2101)
- Access Time from Negative-Going Edge of Memory Read Pulse (\overline{RD}) Until Output Data Valid Must be Less than t_{RDD} (25ns max for ADSP-2101 Operating at 12.5MHz). Otherwise Software Wait States Must be Added, or Processor Clock Frequency Reduced.
- Output Data Must Remain Valid for t_{RDH} (0ns for ADSP-2101)
- Peripheral Device Must Accept Minimum Output Enable Pulse Width of t_{RP} (30ns for ADSP-2101 Operating at 12.5MHz).

Figure 9.3

The t_{RDD} specification determines the peripheral device data access requirement. In the case of the ADSP-2101, $t_{RDD} = 25$ ns minimum. If the access time of the peripheral is greater than this, wait states must be

added or the processor speed reduced. The relationship between these parameters for the ADSP-2101 is given by the following equations:

ADSP-2101 PARALLEL READ TIMING

- t_{ASR} = Address and Data Memory Select Setup Before READ DATA low
- $t_{ASR} = 0.25t_{CK} - 15$ ns Minimum
- t_{RDD} = READ DATA LOW to Data Valid
- t_{CK} = Processor Clock Period (80ns Minimum)
- $t_{RDD} = 0.5t_{CK} - 15$ ns + # wait states * t_{CK} Maximum

Figure 9.4

The ADSP-2101 can easily be interfaced to slow peripheral devices using its programmable wait state generation capability. Three registers control wait state generation for Boot, Program, and Data Memory interfaces. You can specify 0 to 7 wait states for each parallel memory interface. Each wait state added increases the allowable external

data memory access time by an amount equal to the processor clock period (80ns for the ADSP-2101 operating at 12.5MHz). The Data Memory Address, DMS, and RD lines are all held stable for an additional amount of time equal to the duration of the wait states.

PARALLEL ADC TO DSP INTERFACE

The conversion process in a sampling ADC is initiated by a pulse often called the *encode command*, or *start-convert*. The leading (or trailing) edge of this pulse causes the internal ADC sample-and-hold to switch from the sample mode to the hold mode so that the conversion process can take place. Extreme care must be taken in order to insure that this pulse is both jitter-free and noise-free. Any sample-to-sample variation in the occurrence of this edge has the same effect as aperture jitter and will produce a corresponding degradation in the overall ADC signal-to-noise ratio. For this reason, the start-convert signal is usually generated by a stable source external to the DSP processor.

The various timing pulses required to carry out the actual internal conversion process (after receipt of the *convert-start* command) may be generated in several ways depending upon the individual ADC design. In some ADCs, the *convert-start* pulse triggers an internal oscillator or timing chain which in turn controls the conversion. In other ADCs, a user-supplied asynchronous external clock is required.

At some point in time after the convert-start pulse edge, the internal ADC conversion process is completed. In the case of a parallel-output ADC, a single pulse called *data valid*, *data ready*, *read data*, *conversion*

complete, *end-of-conversion*, or *busy/interrupt* is asserted. This pulse is used to drive an *interrupt request* input of the DSP processor as shown in the ADC/DSP parallel interface in Figure 9.5. The DSP then places the address of the ADC on the *data memory address bus* and asserts the *data memory select* line which in turn enables the address decoder. The *chip select* input to the ADC is then asserted along with the *read data* line from the DSP. The *read data* line is connected to the *read* input of the ADC. Asserting the *read* line to the ADC enables the tri-state parallel outputs which are connected to the *data memory data bus*. The DSP then reads the ADC data into an internal register on the rising edge of the *read data* pulse. In order for the circuit shown in Figure 9.5 to operate properly, the timing between the two devices must be made compatible. This will be illustrated by considering a representative example of the ADSP-2101 processor interfaced to the AD7871 ADC.

The AD7871 is a 14 bit, 83kSPS ADC which can operate in either the parallel or serial mode. A functional block diagram of the AD7871 is shown in Figure 9.6. The key interface timing specifications for the two devices are compared in Figure 9.7. Specifications for the ADSP-2101 are given for a clock frequency of 12.5MHz.

ADC/ADSP-2101 PARALLEL INTERFACE

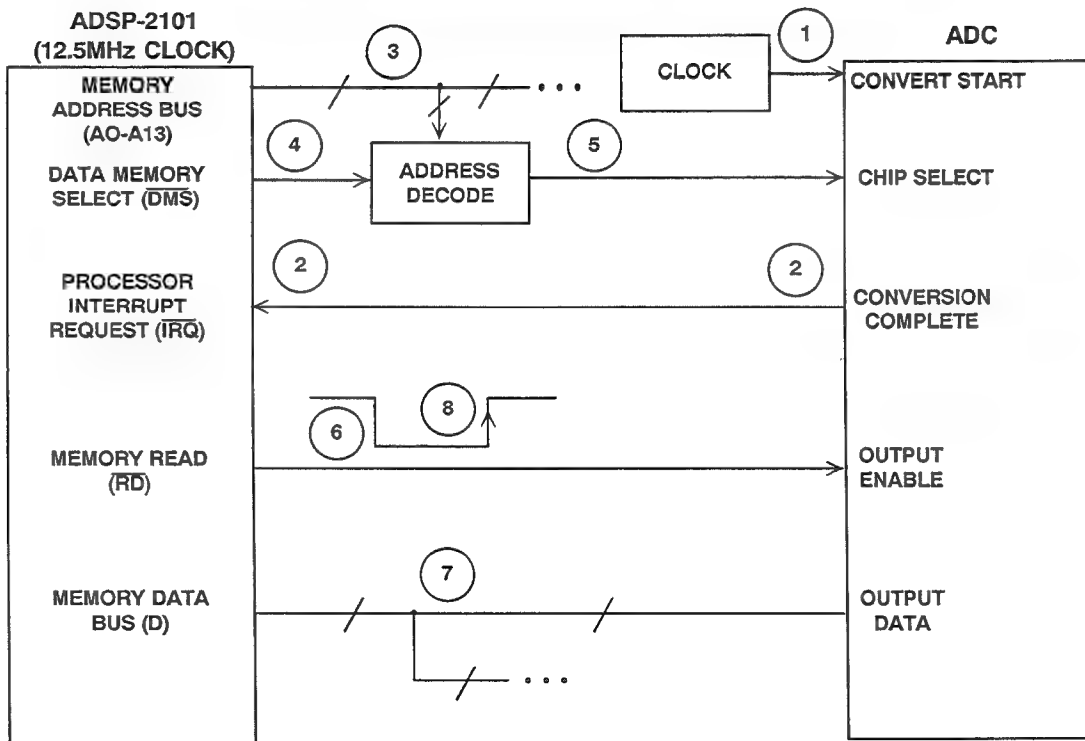


Figure 9.5

AD7871 14-BIT, 83 kSPS ADC FUNCTIONAL DIAGRAM

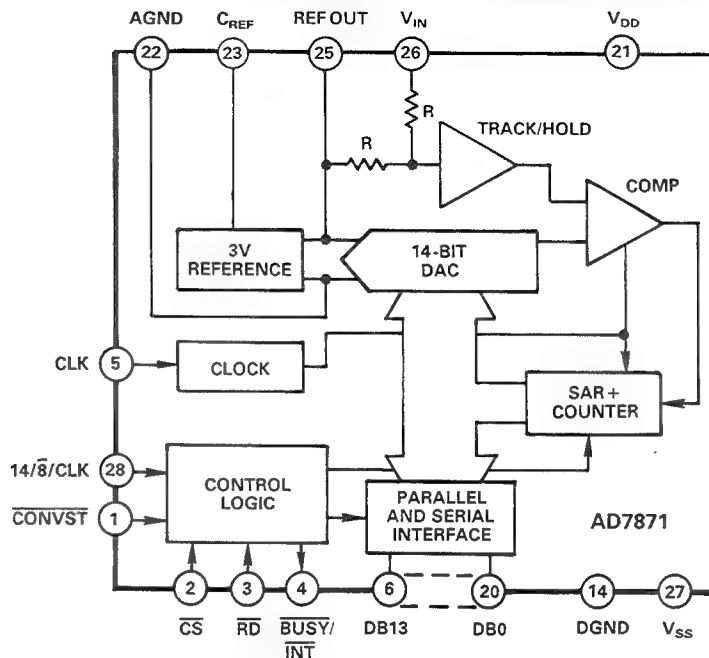


Figure 9.6

ADSP-2101 AND AD7871 PARALLEL *READ* INTERFACE TIMING SPECIFICATIONS

ADSP-2101 PROCESSOR (12.5MHz)	AD7871 ADC
t_{ASR} (Data Address, Data Memory Select Setup Time Before \overline{RD} Low) = 5ns min	t_2 (\overline{CS} to \overline{RD} Setup Time) = 0ns min (Must Add Address Decode Time to This Value)
t_{RP} (\overline{RD} Pulse Width) = 30ns + # wait states * 80ns min	t_3 (\overline{RD} Pulse Width) = 60ns min
t_{RDD} (\overline{RD} Low to Data Valid) = 25ns + # wait states * 80ns min	t_6 (Data Access Time After \overline{RD}) = 57ns max
t_{RDH} (Data Hold from \overline{RD} High) = 0ns min	t_7 (Bus Relinquish Time after \overline{RD}) = 5ns min

Figure 9.7

Examining the timing specifications shown in Figure 9.7 reveals that for the timing between the devices to be compatible, one software wait state must be programmed into the ADSP-2101. A simplified interface diagram for the two devices is shown in Figure 9.8. The conversion complete signal from the AD7871 is designated BUSY/INT.

Parallel interfaces with other DSP processors can be designed in a similar manner by carefully examining the timing specifications for all appropriate signals for each device. The interface between the ADSP-2100 microprocessor (J-Grade, 6.144MHz Clock) is shown in Figure 9.9. Interfacing the AD7871 ADC to faster versions of the ADSP-2100

series requires the addition of wait states using the *Data Memory Acknowledge* (DMACK) signal. The DMACK signal indicates that the memory-mapped peripheral is ready for data transfer. If DMACK is not asserted when checked by the processor, wait states are automatically generated until DMACK is asserted. A detailed description wait state generation in the ADSP-2100 using the DMACK signal and the external hardware required is given in Reference 1 (also included at the end of this section).

The parallel interface between the AD7871 and the TMS32020/C25 is shown in Figure 9.10.

AD7871 ADC PARALLEL INTERFACE TO ADSP-2101




Figure 9.8

AD7871 PARALLEL INTERFACE TO ADSP-2100

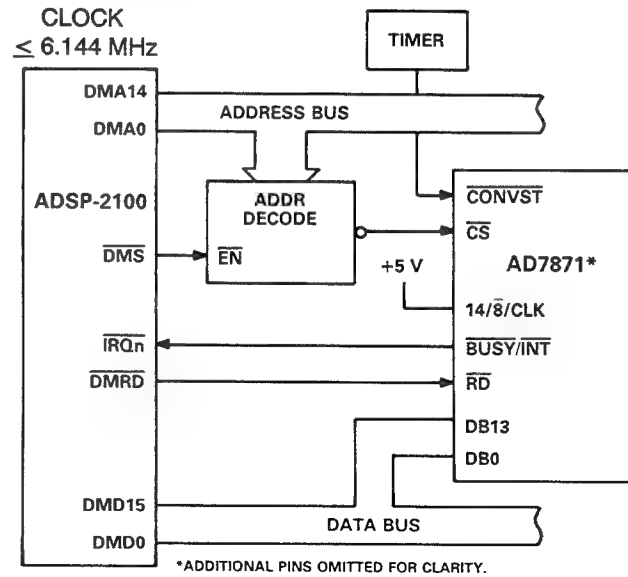


Figure 9.9

AD7871 PARALLEL INTERFACE TO TMS320/C25

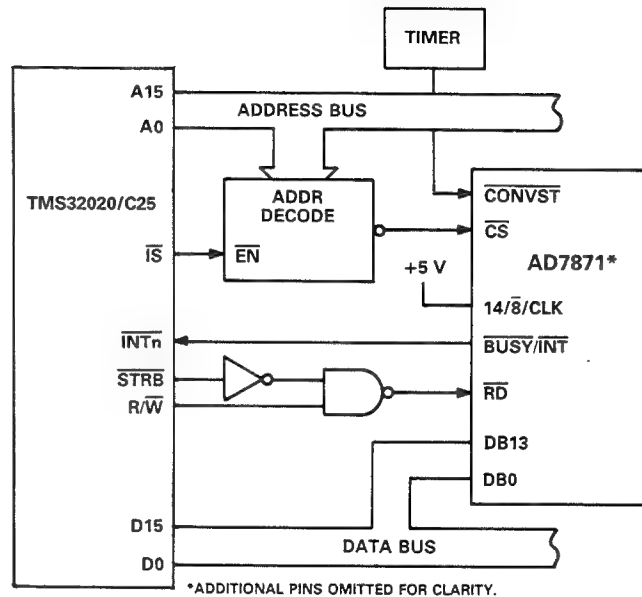


Figure 9.10

PARALLEL INTERFACING TO DSP PROCESSORS: WRITING DATA TO MEMORY-MAPPED DACs

A simplified block diagram of a typical DSP interface to a peripheral device showing write-mode signals is shown in Figure 9.11. The memory-write cycle timing diagram for the ADSP-2101 is shown in Figure 9.12. The write process may be initiated by the peripheral device by asserting the DSP *interrupt request* line indicating that the peripheral is ready to accept a new parallel data word. The DSP then places the address of the peripheral device on the *address bus* and asserts the *data memory select* (DMS) line. This causes the output of the address decoder to assert the *chip select* input to the

peripheral. The *write* (\overline{WR}) output of the DSP is asserted t_{ASW} ns after the negative-going edge of the \overline{DMS} signal. The width of the \overline{WR} pulse is t_{WP} ns. Data is placed on the data bus (D) and is valid t_{Dw} ns before the \overline{WR} line goes high. The positive-going transition of the \overline{WR} line is used to clock the data on the data bus (D) into the external parallel memory. The data on the data bus remains valid for t_{DH} ns after the positive-going edge of the \overline{WR} signal. The key timing requirements for the peripheral device are shown in Figure 9.13.

DAC/ADSP-2101 PARALLEL INTERFACE

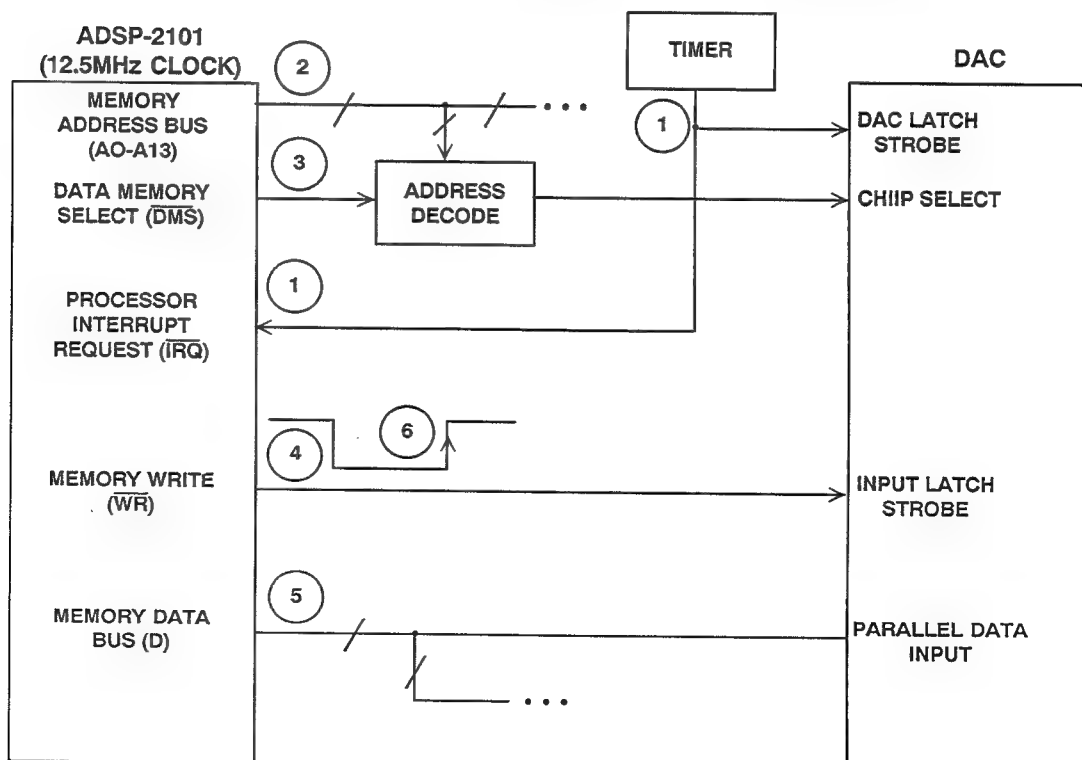


Figure 9.11

ADSP-2101 MEMORY WRITE TIMING

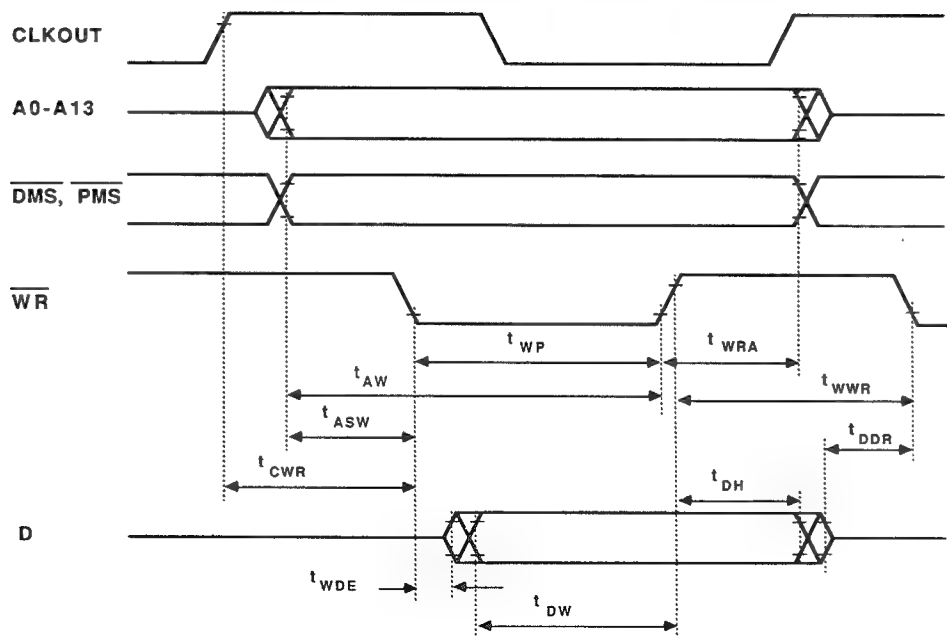


Figure 9.12

PARALLEL PERIPHERAL DEVICE WRITE INTERFACE KEY REQUIREMENTS

- Address Decode Time Plus Peripheral Chip Select Setup Time Must be Less Than Address and Data Memory Select Setup Time t_{ASW} (5ns for ADSP-2101 Operating at 12.5MHz)
- Input Data Setup Time Must be Less Than t_{DW} (20ns for ADSP-2101 Operating at 12.5MHz)
- Input Data Hold Time Must be Less Than t_{DH} (10ns for ADSP-2101 Operating at 12.5MHz)
- Peripheral Device Must Accept Input Write Clock Pulse of Width t_{WP} (30ns min for ADSP-2101 Operating at 12.5MHz)

Figure 9.13

If any of the timing constraints shown in Figure 9.13 are violated by the peripheral device, wait states must be added or the

processor speed reduced. The relationship between these parameters for the ADSP-2101 are shown in Figure 9.14.

ADSP-2101 PARALLEL WRITE TIMING

- t_{CK} = Processor Clock Period (80ns Minimum)
- t_{ASW} = Address and Data Memory Select Time Before \overline{WR} Low = $0.25t_{CK}$ - 15ns Minimum
- t_{DW} = Data Setup Before \overline{WR} High = $0.5t_{CK}$ - 20ns + #Wait States * t_{CK}
- t_{DH} = Data Hold After \overline{WR} High = $0.25t_{CK}$ - 10ns
- t_{WP} = \overline{WR} Pulse Width = $0.5t_{CK}$ - 10ns + #Wait States * t_{CK}

Figure 9.14

The ADSP-2101 can easily be interfaced to slow peripheral devices using its program-mable wait state generation capability which causes the Memory Address, \overline{DMS} , \overline{WR} , and

Data Output lines to remain stable for an amount of additional time equal to the duration of the wait states.

PARALLEL DAC TO DSP INTERFACE

A typical parallel interface between a DSP and a DAC is shown in Figure 9.15. In most DSP applications the DAC is operated continuously from a stable clock source which is external to the DSP processor. The DAC should have double buffering: an input latch to handle the asynchronous DSP interface, and a second latch which drives the DAC current switches. The DAC latch strobe is derived from the external stable clock. In addition to clocking the DAC latch, the DAC latch strobe is also used to generate a processor interrupt which indicates the DAC is ready for new input data. The processor then asserts the *data memory select* line and places the DAC address on the memory address bus. The DAC *chip select* is then asserted, and the *data memory write* line loads the next data word on the data memory

data bus into the DAC input latch. This completes the write cycle, and the DAC is now ready to receive the next DAC latch strobe from the external source. In order for the circuit shown in Figure 9.15 to operate properly, the timing between the two devices must be made compatible. This will be illustrated by considering a representative example of the ADSP-2101 processor interfaced to the AD7840 DAC.

The AD7840 is a 14 bit 100kSPS DAC which has both parallel and serial interface capability. A block diagram of the device is shown in Figure 9.16. The key interface timing specifications for the two devices are compared in Figure 9.17. Specifications for the ADSP-2101 are given for a clock frequency of 12.5MHz.

DAC/ADSP-2101 PARALLEL INTERFACE

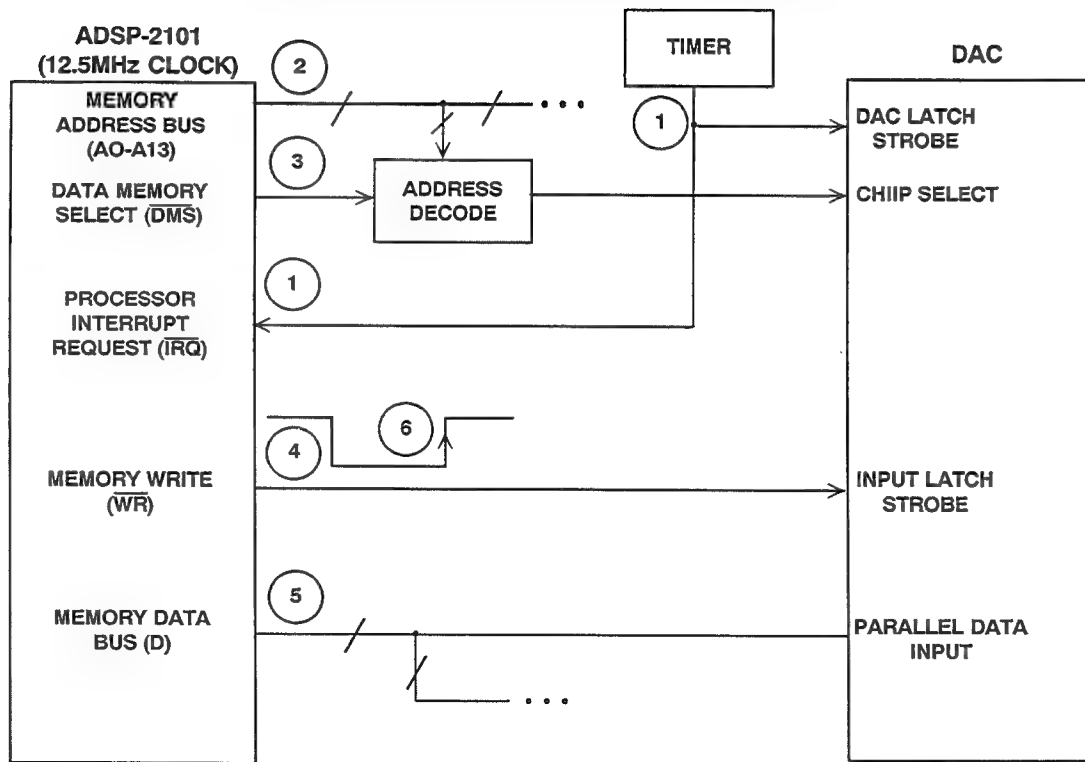


Figure 9.15

AD7840 14-BIT, 100 kSPS DAC FUNCTIONAL DIAGRAM

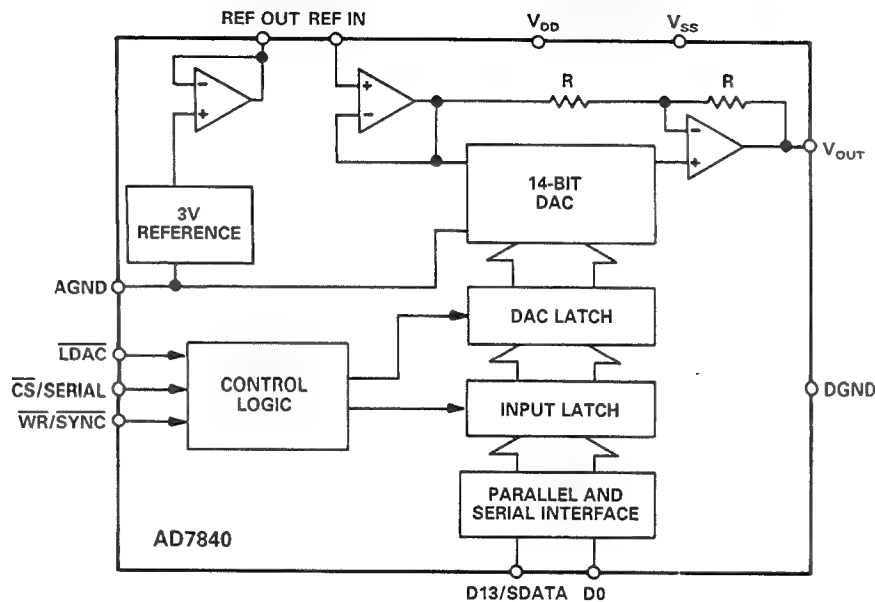


Figure 9.16

ADSP-2101 AND AD7840 PARALLEL *WRITE* INTERFACE TIMING SPECIFICATIONS

ADSP-2101 PROCESSOR (12.5MHz)	AD7840 DAC
t_{ASW} (Address and Data Memory Select Setup Before \overline{WR} Low = 5ns min	$t_1 = \overline{CS}$ to \overline{WR} Setup Time = 0ns min (Must Add Address Decode Time)
t_{WP} (\overline{WR} Pulse Width) = 30ns + # wait states * 80ns min	t_3 (\overline{WR} Pulse Width) = 45ns min
t_{DW} (Data Setup Before \overline{WR} High) = 20ns + # wait states * 80ns min	t_4 (Data Valid to \overline{WR} Setup Time) = 21ns min
t_{DH} (Data Hold After \overline{WR} High) = 10ns	t_5 (Data Valid to \overline{WR} Hold Time) = 10ns min

Figure 9.17

Examining the timing specifications shown in Figure 9.17 reveals that for the timing between the devices to be compatible, at least one software wait state must be programmed into the ADSP-2101. A simplified interface diagram for the two devices is shown in Figure 9.18.

Parallel interfaces with other DSP processors can be designed in a similar manner by carefully examining the timing specifications for all appropriate signals for each device.

The interface between the ADSP-2100 microprocessor (clock speeds up to 8.192MHz) is shown in Figure 9.19. Interfacing the AD7840 to the ADSP-2100A at clock speeds of greater than 8.144MHz requires the addition of wait states using the ADSP-2100 DMACK signal as described in Reference 1.

The parallel interface between the AD7840 DAC and the TMS32020/C25 is shown in Figure 9.20.

AD7840 DAC PARALLEL INTERFACE TO ADSP-2101




Figure 9.18

AD7840 PARALLEL INTERFACE TO ADSP-2100

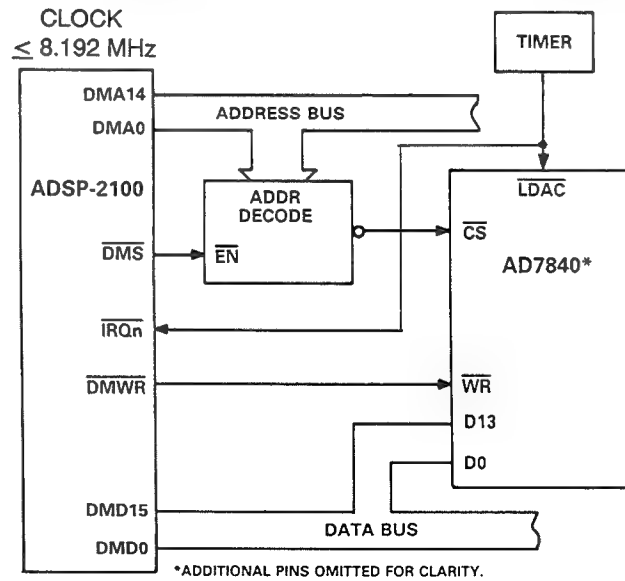


Figure 9.19

AD7840 PARALLEL INTERFACE TO TMS32020

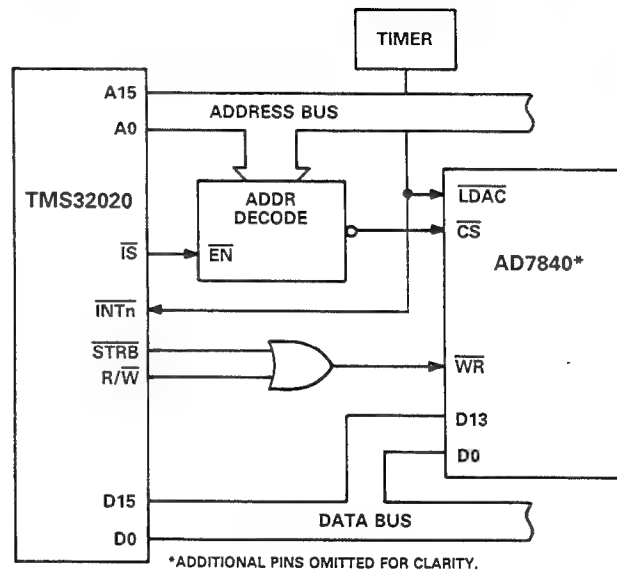


Figure 9.20

SERIAL INTERFACING TO DSP PROCESSORS

DSP processors which have serial ports (such as the ADSP-2101, DSP56000, and the TMS32020/C25) provide a simple interface to peripheral ADCs and DACs. Use of the serial port eliminates the need for using large parallel buses to connect the ADCs and DACs to the DSP. In order to understand serial data transfer better, we will first examine the serial port operation of the ADSP-2101.

A block diagram of one of the two serial ports of the ADSP-2101 is shown in Figure 9.21. The *Transmit* (TX) and *Receive* (RX) registers are identified by name in the ADSP-2101 assembly language, not memory mapped.

In the receiving portion of the serial port, the *Receive Frame Synchronization* (RFS) signal initiates reception. The serial *Receive Data* (DR) from the external device (ADC) is transferred into the *Receive Shift Register* one bit at a time. The negative-going edge of the *Serial Clock* (SCLK) is used to clock the serial data from the external device into the *Receive Shift Register*. When a complete word has been received, it is written to the *Receive Register* (RX), and the receive interrupt for that serial port is generated. The *Receive Register* is then read by the processor.

Writing to the *Transmit Register* readies the serial port for transmission. The *Transmit Frame Synchronization* (TFS) signal initiates transmission. The value in the

Transmit Register (TX) is then written to the internal *Transmit Shift Register*. The data in the *Transmit Shift Register* is sent to the peripheral device (ADC) one bit at a time, and the positive-going edge of the *Serial Clock* (SCLK) is used to clock the serial *Transmit Data* (DT) into the external device. When the first bit has been transferred, the Serial Port generates the transmit interrupt. The *Transmit Register* can then be written with new data, even though the transmission of the previous data is not complete.

In the *normal* framing mode, the frame sync signal (RFS or TFS) is checked at the falling edge of SCLK. If the framing signal is asserted, data is available (transmit mode) or latched (receive mode) on the *next* falling edge of SCLK. The framing signal is not checked again until the word has been transmitted or received. In the *alternate* framing mode, the framing signal is asserted in the *same* SCLK cycle as the first bit of a word. The data bits are latched on the falling edge of SCLK, but the framing signal is checked only on the first bit. Internally-generated framing signals remain asserted for the length of the serial word. The *alternate* framing mode of the serial port in the ADSP-2101 is normally used to receive data from ADCs and transmit data to DACs.

The key features of the ADSP-2101 serial ports are summarized in Figure 9.22.

ADSP-2101 SERIAL PORT BLOCK DIAGRAM

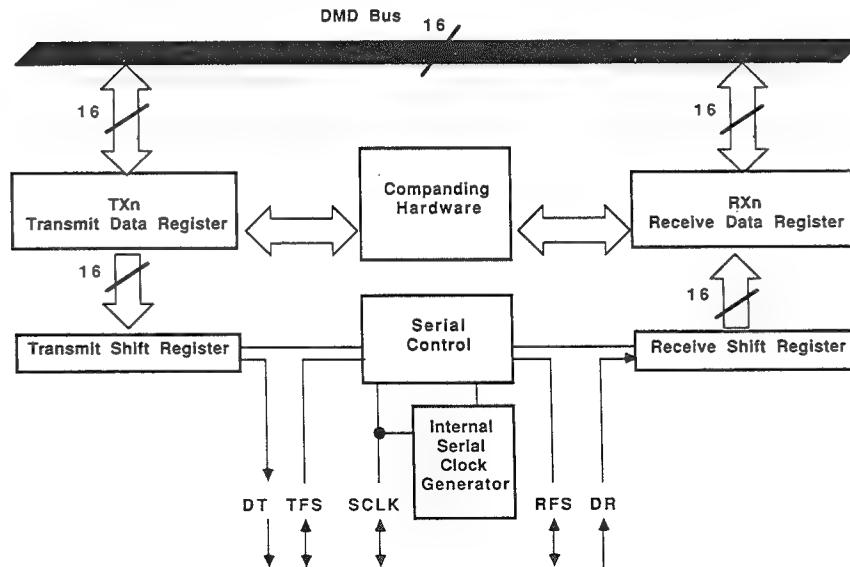


Figure 9.21

ADSP-2101 SERIAL PORTS KEY FEATURES

- Separate Transmit and Receive Sections for Each Port
- Double-Buffered Transmit and Receive Registers
- Serial Clock Can Be Internally (up to 6.25MHz) or Externally (up to 12.5MHz) Generated
- Transmit and Receive Frame Sync Signals Can be Externally or Internally Generated
- Serial Data Words of 3 to 16 Bits Supported
- Automatically Generated Processor Interrupts
- Hardware Companding Capability

Figure 9.22

SERIAL ADC TO DSP INTERFACE

A timing diagram of the ADSP-2101 serial port operating in the receive mode (alternate framing) is shown in Figure 9.23. The first negative-going edge of the SCLK to occur after the rising edge of the RFS input clocks

the MSB data from the ADC into the serial input latch. The process continues until all serial bits have been transferred into the serial input latch. The key timing specifications of concern are the serial data setup

(t_{SCS}) and hold times (t_{SCH}) with respect to the negative-going edge of the SCLK. In the case of the ADSP-2101, these values are both 10ns minimum. The RFS setup and hold times are also 10ns, respectively. Most peripheral ADCs will have no trouble meeting these specifications, even at the maximum serial data transfer rate of 12.5MHz.

The AD7872 ADC is a 14 bit, 83kSPS serial-only version of the AD7871. A block diagram of the device is shown in Figure 9.24. The device operates on a 2MHz external or internal clock. Figure 9.25 shows the AD7872 interfaced to the ADSP-2101. The timing diagram of the AD7872 is shown in Figure 9.26. The SSTRBbar signal is active-low, so the ADSP-2101 must be programmed to accept an inverted RFS input. The serial clock operates at a frequency of 2MHz (500ns period). The serial clock can be programmed for either continuous or gated operation. In this example, it operates in the continuous mode. The data bits are valid t_{12} ns (155ns max) after the positive-going edges of SCLK. This allows a setup time of $250 - 155 = 95$ ns minimum before the negative-going edges of SCLK. The hold-time after the negative-

going edge of SCLK is therefore at least equal to one-half the clock period, or 250ns. The positive-going edge of the SSTRB signal occurs t_{13} ns (140ns max) after the positive-going edge of SCLK after the last data bit is transferred. This allows $250 - 140 = 110$ ns minimum before the next negative-going edge of SCLK. These simple calculations show that the data and RFS setup and hold requirements of the ADSP-2101 (10ns) are met with considerable margin.

The ADSP-2101 can be easily programmed to generate the 2MHz serial clock for the AD7872 if desired. Details can be found in the ADSP-2101 User's Manual/Architecture. The Convert Start (CONVST) signal is generated externally to the AD7872 from a stable clock source which is asynchronous to the Serial Clock.

The serial interface between the AD7872 and the DSP56000 is shown in Figure 9.27, and the interface with the TMS32020/C25 is shown in Figure 9.28. The simple interfaces shown to the three DSP processors are referred to as *zero-chip interfaces* because no additional *glue* logic is required.

ADSP-2101 SERIAL PORT RECEIVE TIMING




Figure 9.23

AD7872 ADC SERIAL INTERFACE TIMING

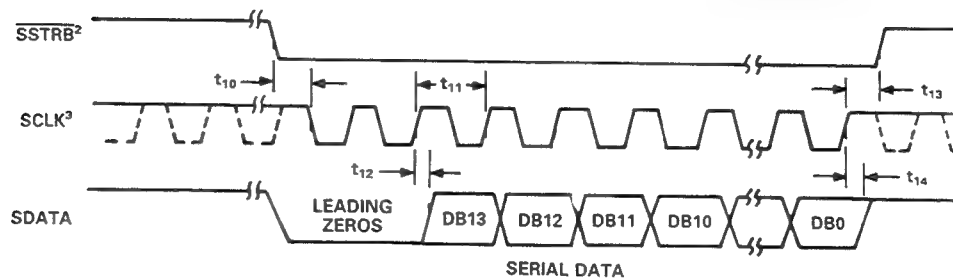


Figure 9.26

AD7872 SERIAL INTERFACE TO DSP56000

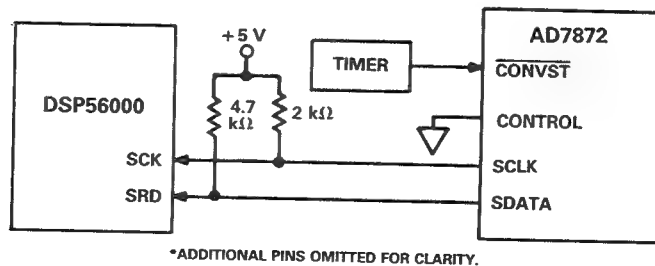


Figure 9.27

AD7872 SERIAL INTERFACE TO TMS32020/C25

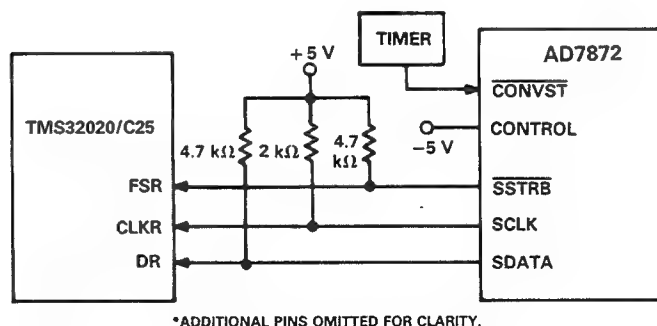


Figure 9.28

SERIAL DAC TO DSP INTERFACE

A timing diagram of the ADSP-2101 serial port operating in the alternate framing transmit mode (with internally generated Transmit Frame Sync) is shown in Figure 9.29. The first negative-going edge of the SCLK to occur after the rising edge of the TFS output clocks the MSB data from the serial port into the DAC serial input latch. The process continues until all serial bits have been transferred into the DAC serial input latch. The key timing specifications of concern are the data output setup and hold times with respect to the negative-going edge of the SCLK. The ADSP-2101 specifies that the TFS output will be a valid high t_{RH} ns (15ns max) after the positive-going edge of SCLK. The serial transmit data is valid t_{SCDV} ns (25ns max) after the positive-going edge of SCLK. Due to the high speed of the serial port interface of the ADSP-2101, data setup and hold times are therefore approximately equal to one-half the period of the serial clock for clock rates up to 12.5MHz.

The AD766 is a 16 bit serial DAC which can operate at sample rates up to 500kSPS and is fully specified in terms of both dc and ac parameters such as THD and SNR. A block diagram of the device is shown in Figure 9.30. Data is transmitted to the AD766 in a bit stream composed of 16 bit words with a serial, MSB first format. Three signals must be present to achieve proper operation: the *data*, *clock*, and *latch enable* signals. Input data bits are clocked into the input register on the falling edge of the clock signal. The LSB is clocked in on the 16th clock pulse. When all data bits are loaded, a low-going *latch enable* pulse updates the DAC input. Figure 9.31 illustrates the general signal requirements for data transfer for the AD766. Data setup and hold-times (with respect to the negative-going SCLK edge) are each 15ns. The negative-going edge of the latch enable must occur at least 15ns before the negative-going edge of SCLK. These detailed timing requirements are

ADSP-2101 SERIAL PORT TRANSMIT TIMING

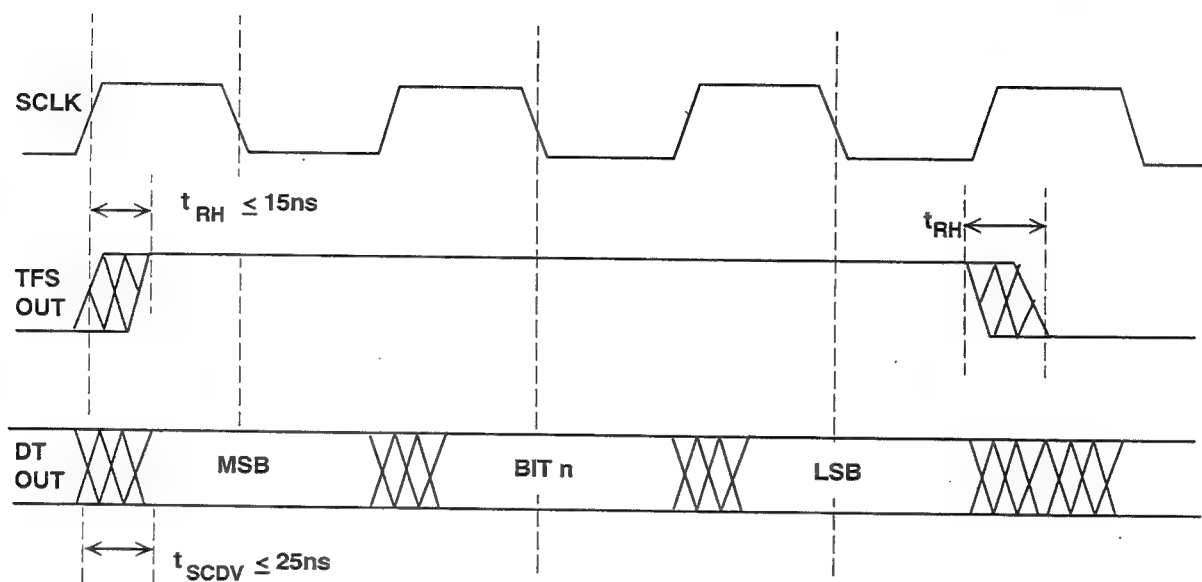


Figure 9.29

AD766 16-BIT, 500kSPS DSP DAC FUNCTIONAL DIAGRAM

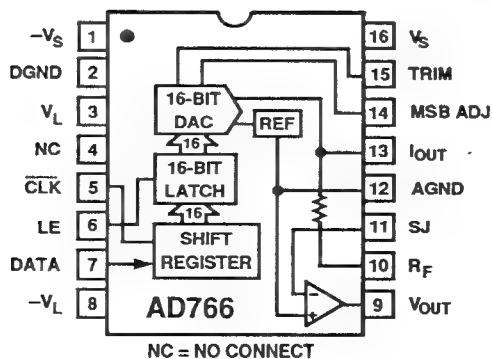


Figure 9.30

AD766 DAC SIGNAL REQUIREMENTS




Figure 9.31

AD766 TIMING REQUIREMENTS




Figure 9.32

illustrated in Figure 9.32. These timing requirements are compatible with the serial ports of popular DSP processors. The AD766 input clock can run at a 12.5MHz rate. This clock rate will allow sampling rates up to 500kSPS.

The ADSP-2101 incorporates two complete serial ports which can be directly interfaced

to the AD766 as shown in Figure 9.33. Using both serial ports, two AD766's can be directly interfaced with no additional hardware. The zero-chip interface to the TMS320C25 is shown in Figure 9.34. The maximum serial clock rate for the TMS320C25 is 5MHz. Figure 9.35 shows the serial interface to the DSP56000/56001.

AD766 SERIAL INTERFACE TO ADSP-2101

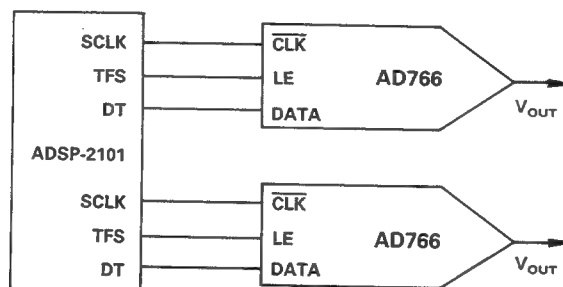


Figure 9.33

AD766 SERIAL INTERFACE TO TMS320C25




Figure 9.34

AD766 SERIAL INTERFACE TO DSP56000




Figure 9.35

INTERFACING I/O PORTS AND CODECS TO DSPs

Since most DSP applications require both an ADC and a DAC, I/O Ports and Codecs have been developed which integrate the two functions on a single chip as well as provide easy-to-use interfaces to standard DSPs.

A functional block diagram of the AD7868 12 bit, 83kSPS I/O Port is shown in Figure 9.36. The AD7869 is a 14 bit I/O Port which is functionally equivalent to the AD7868. These devices are fully specified in terms of ac and dc performance. The SNR (including distortion) of the AD7868 is 72dB, while the

AD7869 is 82dB. Both devices provide simple interfaces to the serial ports of standard DSP microcomputers such as the ADSP-2101 (see Figure 9.37), TMS3020/C25, and the DSP56000.

The ADSP-28msp02 is a complete voiceband codec (ADC and DAC) based on sigma-delta technology. A block diagram is shown in Figure 9.38. The device provides a complete analog front end for high performance voiceband DSP applications. Key features of the device are given in Figure 9.39.

AD7868 12-BIT, 83kSPS I/O PORT

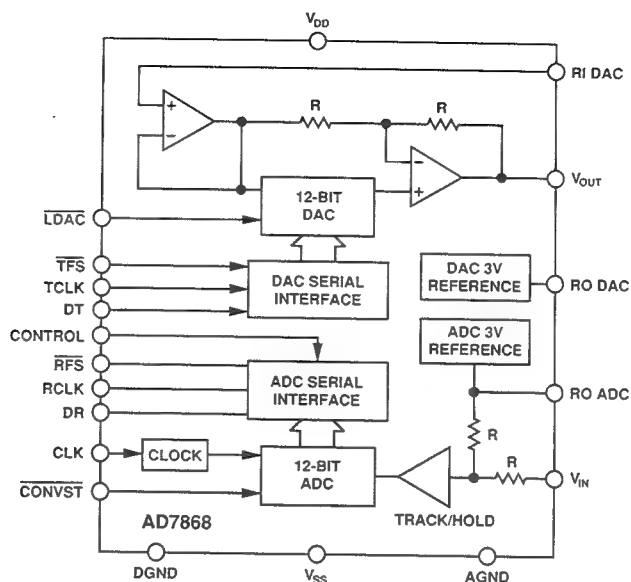
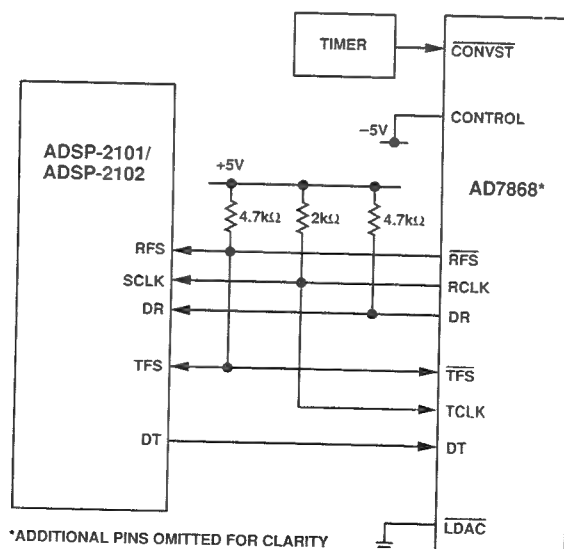


Figure 9.36

AD7868 I/O PORT INTERFACE TO ADSP-2101



*ADDITIONAL PINS OMITTED FOR CLARITY

Figure 9.37

ADSP-28msp02 SIGMA-DELTA VOICEBAND CODEC

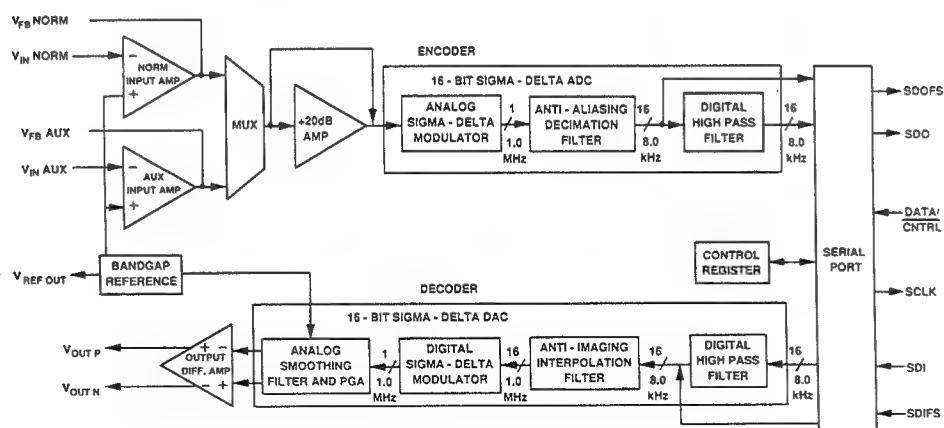


Figure 9.38

KEY FEATURES OF THE ADSP-28msp02 SIGMA-DELTA CODEC

- 16 bit 128x Oversampling Sigma-Delta ADC and DAC
- On-Chip Antialiasing and Smoothing Filters
- On-Chip Input and Output PGAs
- 8kSPS Sampling Rate
- 65dB SNR and THD
- Easy to Interface to Serial Port of Standard DSP Chips
- 24-pin DIP/SOIC Package
- Single +5V Supply, 100mW Power Dissipation
- Ideal for Voiceband Applications

Figure 9.39

Compared to traditional m-law and A-law codecs, the ADSP-28msp02's linear coded ADC and DAC maintain wide dynamic range throughout the transfer function. The encoder side of the device consists of two selectable analog input amplifiers and a sigma-delta ADC. The gain of the input amplifiers can be adjusted with the use of external resistors from -12dB to +26dB. A optional 20dB preamplifier can be inserted before the ADC. The preamplifier and the multiplexer

are configured by bits in the control register. The decoder consists of a sigma-delta DAC and a differential amplifier. The output of the DAC drives an analog smoothing filter which converts the data into an analog voltage. The gain of the smoothing filter and PGA can be adjusted via the control register from -15dB to +6dB in 3dB steps. The ADSP-28msp02 easily interfaces to the serial ports of popular DSP microcomputers such as the ADSP-2101 as shown in Figure 9.40.

ADSP-28msp02 CODEC SERIAL PORT DSP INTERFACE




Figure 9.40

SERIAL VERSUS PARALLEL DSP INTERFACE SUMMARY

Some DSP processors such as the ADSP-2100 support only memory-mapped peripherals and have no serial port. A large number of peripheral devices can be connected to the parallel address and data bus. Each device is treated as a single location in the data memory. A number of high performance ADCs and DACs are available with parallel interfaces. Data setup and hold specifications, write and read pulse widths, etc., must be examined carefully to insure that there are no interface timing violations. Conflicts frequently occur because DSP processors are designed to operate at clock frequencies often

exceeding 10MHz, while ADCs and DACs used in most DSP applications rarely exceed sampling rates of 500kSPS. These interfacing timing conflicts can usually be resolved with the addition of software or hardware wait states.

ADCs which interface to parallel DSPs must have tri-state outputs so that the data bus can be shared among other peripherals. The convert-start signal for the ADC is generated externally to minimize jitter. The conversion-complete signal is typically used to generate an interrupt request to the DSP processor. Care must be taken in the routing

of the ADC parallel digital outputs to prevent digital switching noise from coupling into the ADC analog input.

DACs which interface to parallel DSPs must have double-buffered digital inputs. The latch which drives the DAC switches is updated continuously with a stable external clock source. The external clock is also used to generate an interrupt request to the processor. The input latch is loaded asynchronously by the DSP processor. Care must be taken in the routing of the digital input signals so that they don't couple into the DAC analog output.

The serial ports provided on modern DSP processors provide several advantages when interfacing to peripheral devices. The inter-

face is simple (three wires) and requires little or no external *glue* logic components. Pin counts are minimized as well as logic switching noise. The serial port handles data transmission, reception, and also generates processor interrupt requests automatically. Serial clock and frame synchronizing signals may be generated either internally or externally.

Serial ADCs and DACs are available which are compatible with the timing of most DSP serial ports, and timing conflicts are rare. The Sigma-Delta and Successive Approximation ADC architecture are popular in DSP applications, and also have a serial output data format.

PARALLEL DSP INTERFACE CHARACTERISTICS

- **Peripherals are Memory-Mapped**
- **Timing Conflicts May Require Software or Hardware Wait States**
- **ADCs Must Have Tri-State Outputs**
- **DACs Must Have Double-Buffered Inputs**
- **Routing of Digital Signals is Critical for Low Noise**

Figure 9.41

SERIAL DSP INTERFACE CHARACTERISTICS

- **Three-Wire, Zero-Chip Interface Typical**
- **Data Transmission, Reception, Processor Interrupts Handled by Serial Port**
- **Serial Clock and Frame Synchronization Signals Generated Internally or Externally**
- **Sigma-Delta and Successive Approximation ADCs are Naturally Serial Output Devices**

Figure 9.42

REFERENCES

1. Kapriel Karagozyan, *Wait State Generation on the ADSP-2100 and the ADSP-2100A*, Analog Devices Application Note E1317-8-8/89. (Included at end of this section)
2. *ADSP-2100 User's Manual/Architecture*, Analog Devices
3. *ADSP-2101/2102 User's Manual/Architecture*, Analog Devices



ONE TECHNOLOGY WAY • P. O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

APPLICATION NOTE

Wait State Generation on the ADSP-2100 and ADSP-2100A

by Kapriel Karagozyan

INTRODUCTION

The ADSP-2100 microprocessor can interface to a large number of peripherals (A/D and D/A converters, FIFOs, etc.) by mapping them onto its data memory address space. Some of these peripherals may not be able to be accessed within a single processor cycle. The data memory interface may require the insertion of hardware wait states during read and/or write operations to such peripherals. The data memory acknowledge (DMACK) input is provided on the ADSP-2100 for this interface.

DMACK TIMING

The ADSP-2100 checks the status of the DMACK signal towards the end of each data memory access cycle. If the DMACK signal is not asserted high, the processor extends the current cycle by another full cycle. This extension occurs until the DMACK signal is sampled high, in which case the access is completed. When no wait states are needed, it is recommended that the DMACK input be tied to a logic 1 (always high).

All of the processor address and control lines are held steady during an extended memory access. The only active output is CLKOUT. The DMACK input is not checked during instruction cycles that do not access the data memory.

DMACK should be held high during normal processor operation. If wait states are desired, external wait state generation logic must return the DMACK input valid low within a required time range to ensure that the DMACK low level is recognized by the processor, which in turn extends the data memory access. The timing requirements relevant to the generation of DMACK are provided with the data memory read and write specifications in the *ADSP-2100 Data Sheet*. They are: DMA valid to DMACK valid (#75), $\overline{\text{DMRD}}$ low to DMACK valid (#74), $\overline{\text{DMWR}}$ low to DMACK valid (#99) and CLKOUT high to DMACK invalid (#103).

The DMACK input to the processor is internally sampled shortly before the rising edge of CLKOUT. The minimum requirement of specification #103 forces the wait state

generation logic to keep the valid DMACK level until the rising edge of CLKOUT (which occurs once towards the end of every processor cycle) and ensures that the DMACK input is acknowledged by the processor. On the other hand, the maximum specification on #103 ensures that the low level DMACK does not run into the next cycle and cause another wait state. In order to complete the data memory access, DMACK must be brought high within the maximum specification in #103 and kept high until the next rising edge of CLKOUT. This ensures that DMACK is high when it is sampled again on the next cycle and thus causes the extended access to finish. To generate any number of wait states, the wait state generation logic should use CLKOUT as a counter to determine when the appropriate number of wait states have elapsed. The requirement #103 implies that CLKOUT must be used to clear the wait state generator, since it is the only actively switching output during an extended memory access.

WAIT STATE GENERATION CIRCUITS

The circuits illustrated in Figures 1, 2 and 3 are recommended implementations of external wait state generation circuitry for the ADSP-2100. The logic components are shown without part numbers because their speed requirements vary with processor speed grades. It is also possible to implement most of the logic using high speed PALs. The circuit shown in Figure 1 extends the memory access by one cycle, whereas the one in Figure 2 extends by two cycles. The circuit in Figure 3 shows how multiple external peripherals and their wait state signals can be gated to generate DMACK.

The circuit in Figure 1 normally keeps the DMACK signal at a logic high until an address requiring wait states is issued by the processor. The output of the D-flop is normally low due to the high level of $\overline{\text{DMWR}}$ or $\overline{\text{DMRD}}$ or the existence of an invalid address. When an address requiring wait states is decoded, DMACK is pulled low. The rising edge of CLKOUT causes the output of the D-flop and consequently the DMACK signal to rise high before the next cycle. Thus, a valid low level DMACK signal is generated, held low and brought back high within the processor's timing specifications. The processor samples DMACK low during the first cycle and extends the data

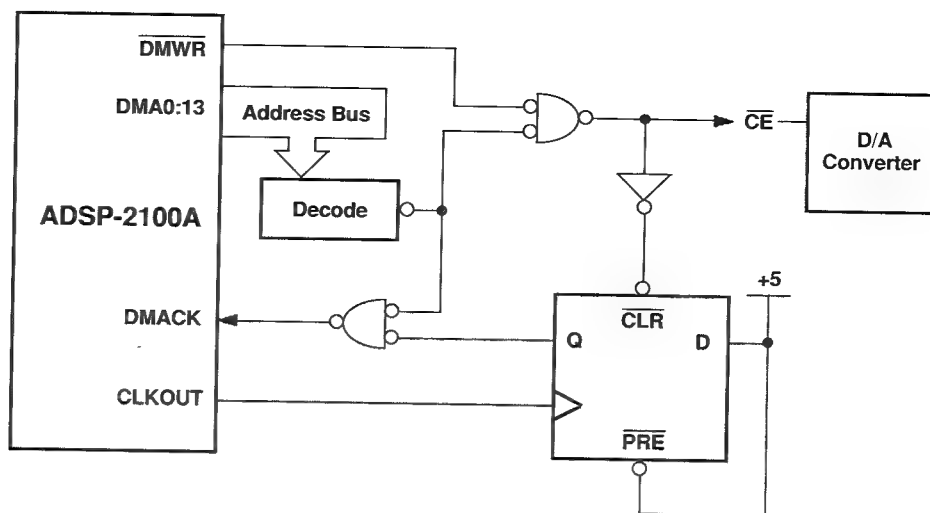


Figure 1. Wait State Generation Logic for One Wait State

memory access. DMACK is high on the second cycle, causing the processor to finish the memory access and resume its normal operation. The DMACK line may have glitches in this circuit configuration, but the circuit guarantees by design that DMACK is valid when needed. Glitches occur only whenever DMACK is not being sampled by the processor and thus are not recognized by the processor.

The operation of the circuit shown in Figure 2 is very similar to that of the one in Figure 1, with the difference that an extra D-flop is added to the DMACK path. This extra D-flop allows the circuit to bring the DMACK signal high after the second rising edge of CLKOUT, hence generating two wait states. More wait states can be generated by adding more D-flops to the DMACK path.

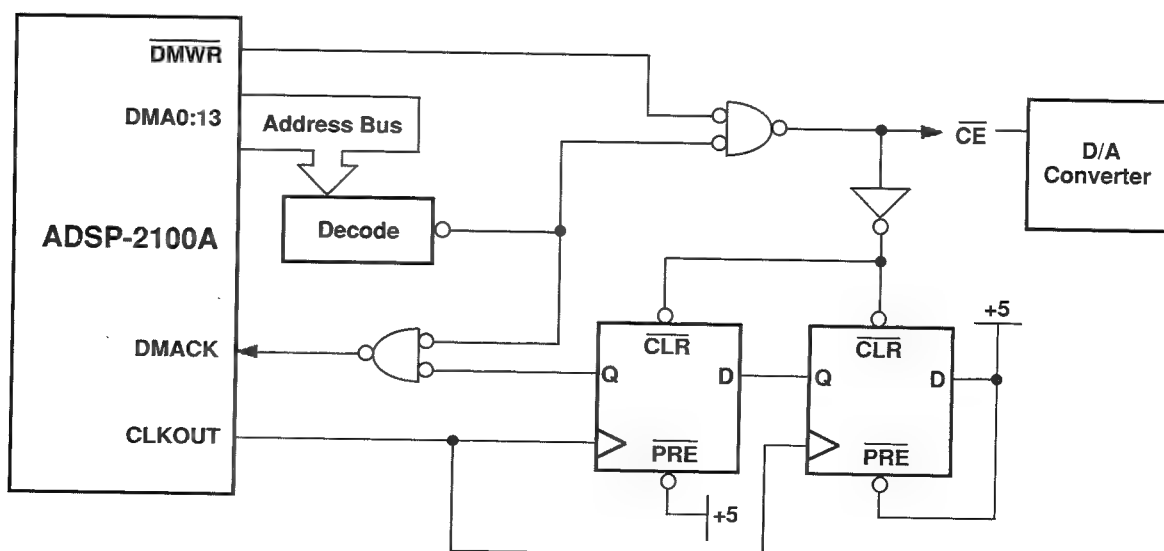


Figure 2. Wait State Generation Logic for Two Wait States

Figure 3 shows how peripherals requiring different numbers of wait states each have their own wait state generation logic. In

this example, reading the A/D converter requires one wait state and writing the D/A converter requires two.

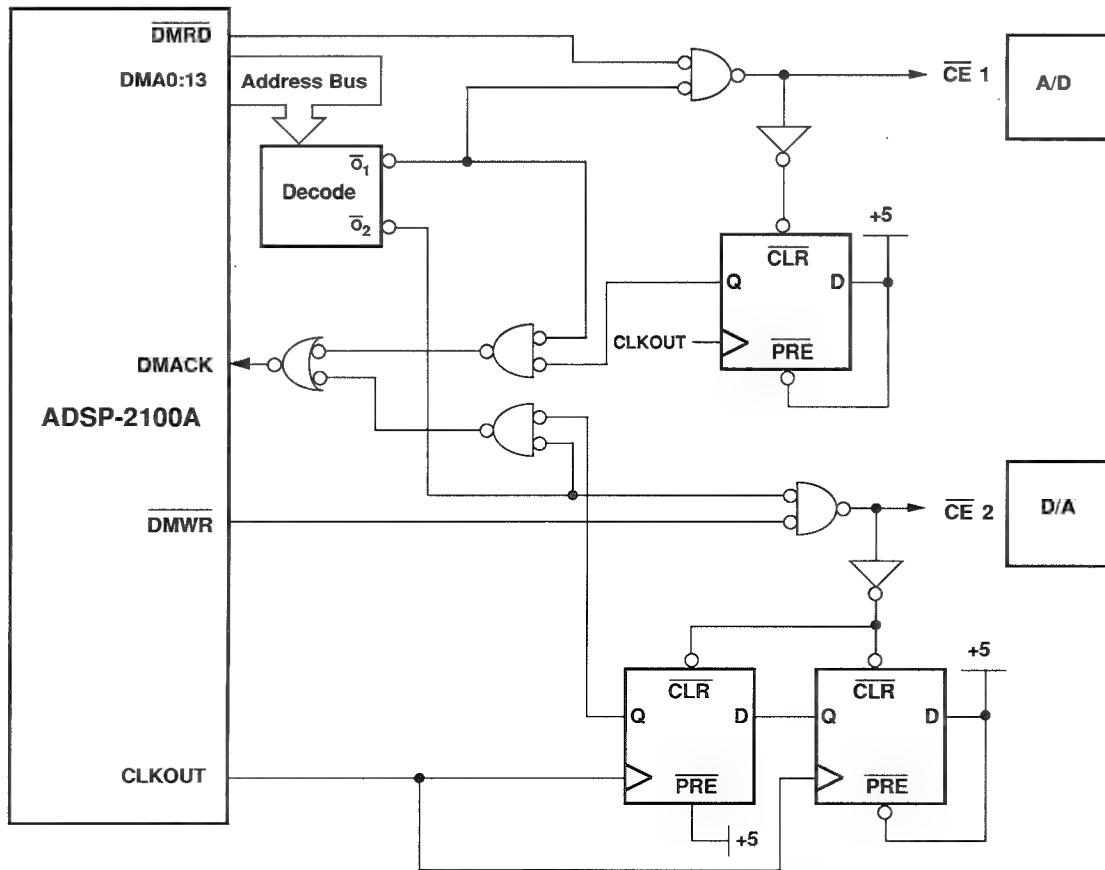
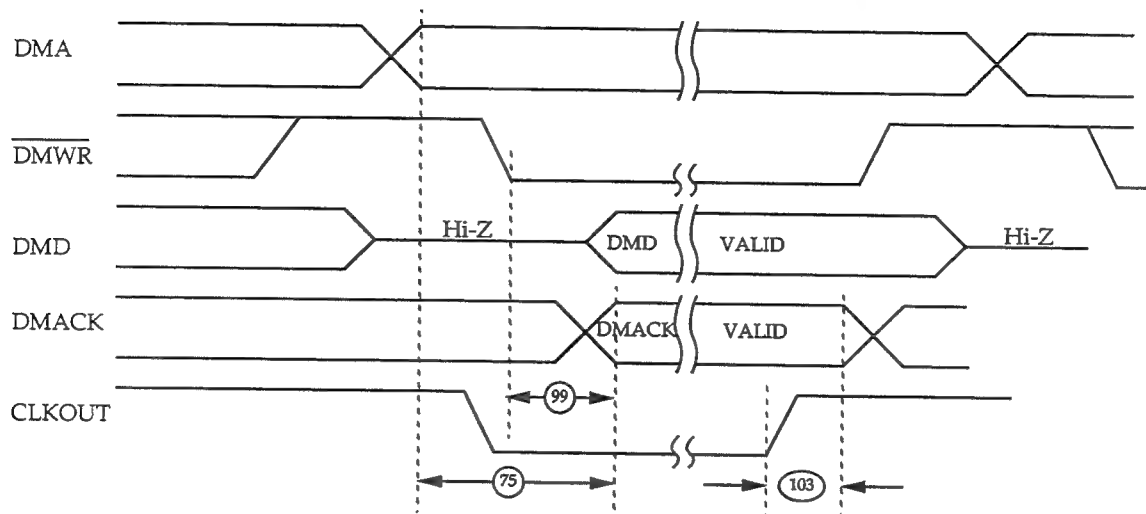


Figure 3. Wait State Generation Logic for Multiple Peripherals

ADSP-2100 DATA MEMORY WRITE CYCLE WITH WAIT STATES



Timing Requirements

- 74 $\overline{\text{DMRD}}$ low to DMACK valid
- 75 DMA valid to DMACK valid
- 99 $\overline{\text{DMWR}}$ low to DMACK valid
- 103 CLKOUT high to DMACK invalid

SECTION X

MIXED SIGNAL PROCESSING APPLICATIONS



MIXED SIGNAL PROCESSING APPLICATIONS

■ HIGH PERFORMANCE MODEMS

V.32 Modem Overview

V.32 Modem Transmitter

V.32 Modem Receiver

I/O Ports and Codecs for V.32 Modems

■ DIGITAL MOBILE RADIO

Overview

The GSM System

Speech Codec

Discontinuous Transmission (DTX)

GSM System Upconversion and Downconversion

■ DIGITAL AUDIO STUDIO RECORDING

■ COMPACT DISC (CD) PLAYER ELECTRONICS

SECTION X

MIXED SIGNAL PROCESSING APPLICATIONS

HIGH PERFORMANCE MODEMS

Modems (*Modulator/Demodulator*) are widely used to transmit and receive digital data using analog modulation over the General Switched Telephone Network (GSTN) as well as private lines. Although the data to be transmitted is digital, the telephone channel is designed to carry voice signals having a bandwidth of approximately 300 to 3300Hz. The telephone transmission channel suffers from delay distortion, noise, crosstalk, near-end and far-end echoes, and other imperfections listed in Figure 10.1. While certain levels of these signal degradations are perfectly acceptable for voice communication,

they can cause high error rates in digital data transmission. The fundamental purpose of the transmitter portion of the modem is to prepare the digital data for transmission over the analog voice line. The purpose of the receiver portion of the modem is to receive the signal which contains the analog representation of the data, and reconstruct the original digital data at an acceptable error rate. High performance modems make use of digital techniques to perform such functions as modulation, demodulation, error detection and correction, equalization, and echo cancellation.

IMPERFECTIONS IN THE TELEPHONE CHANNEL

- Attenuation
- Bandwidth Flatness
- Harmonic Distortion
- Echoes (Near-End and Far-End)
- Phase Jitter
- Phase Distortion, Group Delay Variation
- Noise
- Impedance Mismatches
- Frequency Offset
- Phase and Gain Hits

Figure 10.1

A block diagram of a telephone channel is shown in Figure 10.2. Most voiceband telephone connections involve several connections through the telephone network. The 2-

wire subscriber line available at most sites is generally converted to a 4-wire signal at the telephone central office. The signal is converted back to a 2-wire signal at the far-end

TELEPHONE CHANNEL BLOCK DIAGRAM

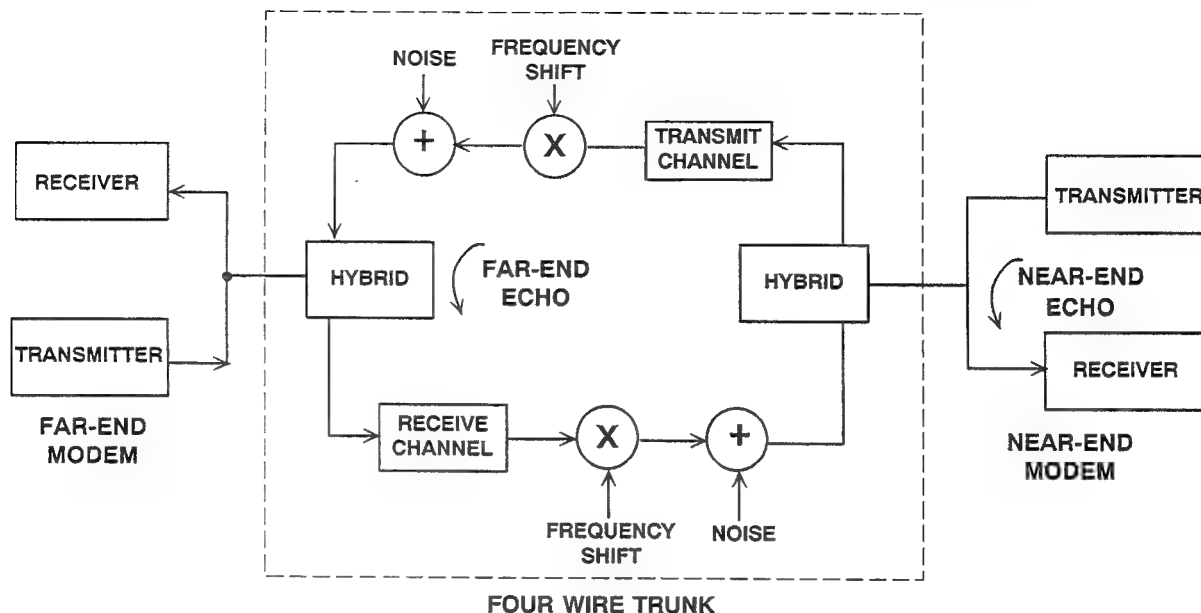


Figure 10.2

subscriber line. The 2- to 4-wire interface is implemented with a circuit called a *hybrid*. The hybrid intentionally inserts impedance mismatches to prevent oscillations on the 4-wire trunk line. The mismatch forces a portion of the transmitted signal to be reflected or echoed back to the transmitter. This echo can corrupt data the transmitter receives from the far-end modem.

Half-duplex modems are capable of passing signals in either direction on a 2-wire line, but not simultaneously. *Full-duplex* modems operate on a 2-wire line and can transmit and receive data simultaneously. Full-duplex operation requires the ability to separate a receive signal from the reflection (echo) of the transmitted signal. This is accomplished by assigning the signals in the two directions different frequency bands separated by filtering, or by echo cancelling in which a locally synthesized replica of the reflected transmitted signal is subtracted from the composite receive signal.

There are two types of echo in a typical voiceband telephone connection. The first echo is the reflection from the near-end hybrid, and the second echo is from the far-end hybrid.

In long distance telephone transmissions, the transmitted signal is heterodyned to and from a carrier frequency. Since local oscillators in the network are not exactly matched, the carrier frequency of the far-end echo may be offset from the frequency of the transmitted carrier signal. In modern applications this shift can affect the degree to which the echo signal can be cancelled. It is therefore desirable for the echo canceller to compensate for this frequency offset.

For transmission over the telephone voice network, the digital signal is modulated onto an audio sinewave carrier, producing a modulated tone signal. The frequency of the carrier is chosen to be well within the telephone band. The transmitting modem modulates the audio carrier with the trans-

mit data signal, and the receiving modem demodulates the tone to recover the receive data signal.

The baseband data signal may be used to modulate the amplitude, the frequency, or the phase of the audio carrier, depending on the data rate required. These three types of modulation are known as amplitude shift keying (ASK), frequency shift keying (FSK), and phase shift keying (PSK). In its simplest form the modulated carrier takes on one of two states - that is, one of two amplitudes, one of two frequencies, or one of two phases. The two states represent a logic 0 or a logic 1.

Low- to medium-speed data links usually use FSK up to 1200 bits/s. Multiphase PSK

are used for 2400 bits/s and 4800 bits/s links. PSK utilizes bandwidth more efficiently than FSK but is more costly to implement. ASK is least efficient and is used only for very low speed links (less than 100 bits/s) For 9600 bits/s, a combination of PSK and ASK is used, known as Quadrature Amplitude Modulation (QAM).

Assuming 7-bit ASCII and 4 bits/character overhead (start, parity, and two stop bits), a data transmission rate of 300 bits/s translates to approximately 27 characters/s. This is faster than a person can type but is too slow for transferring large files or for many applications requiring graphics.

MODULATION METHODS FOR MODEMS

- Amplitude Shift Keying (ASK): Up to 100 bits/s
- Frequency Shift Keying (FSK): Up to 1200 bits/s
- Phase Shift Keying (PSK): Up to 4800 bits/s
- Quadrature Amplitude Modulation (QAM): Up to 9600 bits/s

Figure 10.3

The International Telegraph and Telephone Consultative Committee (CCITT in French) has established standards and

recommendations for modems which are given in Figure 10.4.

CCITT RECOMMENDATIONS FOR TELEPHONE MODEMS

CCITT REC.	DATE	SPEED (BITS/s)	HALF DUPLEX/ FULL DUPLEX/ ECHO CANCEL	GSTN/ PRIVATE	MODULATION METHOD
V.21	1964	200	FDX	GSTN	FSK
V.22	1980	1200	FDX	GSTN	PSK
V.22bis	1984	2400	FDX	GSTN	16QAM
V.23	1964	1200	HDX	GSTN	FSK
V.26	1968	2400	HDX	Private	PSK
V.26 bis	1972	2400	HDX	GSTN	PSK
V.26 ter	1984	2400	FDX(EC)	GSTN	PSK
V.27 bis	1976	4800	HDX	Private	8PSK
V.27 ter	1976	4800	HDX	GSTN	8PSK
V.29	1976	9600	HDX	Private	16AM/PM
V.32	1984	9600	FDX(EC)	GSTN	32QAM
V.33	-	14400	HDX	Private	64QAM

Figure 10.4

V.32 MODEM OVERVIEW

The goal in designing high performance modems is to achieve the highest data transfer rate possible over the GSTN and avoid the expense of using dedicated conditioned private telephone lines. The V.32 recommendation describes a full-duplex (simultaneous transmission and reception) synchronous modem that operates on the Public Switched Telephone Network (GSTN). The V.32 modem communicates at a rate of 9600 bits/s utilizing quadrature amplitude modulation (QAM). Four-bit symbols (bauds) modulate a

carrier frequency of 1800Hz with a modulation rate of 2400 bauds/s. The modulation of 4-bit symbols at a rate of 2400 symbols/s yields the 9600 bits/s specification. These 4-bit symbols are transmitted using 32-state trellis-encoded QAM. The trellis encoding provides an extra bit per symbol for forward error correction. This additional bit dramatically increases the noise performance of the modem. Characteristics of the V. 32 modem are summarized in Figure 10.5

V.32 MODEM CHARACTERISTICS

- **9600 bits/second Bit Rate on GSTN**
- **1800Hz Carrier Frequency (Transmit and Receive)**
- **4 Bits/Symbol, 2400Hz Symbol Rate**
- **32-QAM, Trellis Coded, 4 Bit Data + Redundancy Bit**
- **Transmit/Receive Isolation Using Echo Cancellation**
- **Extensive Use of DSP Techniques**

Figure 10.5

A simplified block diagram for a V.32 modem is shown in Figure 10.6. The diagram shows that the bulk of the signal processing is done digitally. Both the trans-

mit and receive portions of the modem subject the digital signals to a number of DSP algorithms which can be efficiently run on modern processors.

V.32 MODEM BLOCK DIAGRAM

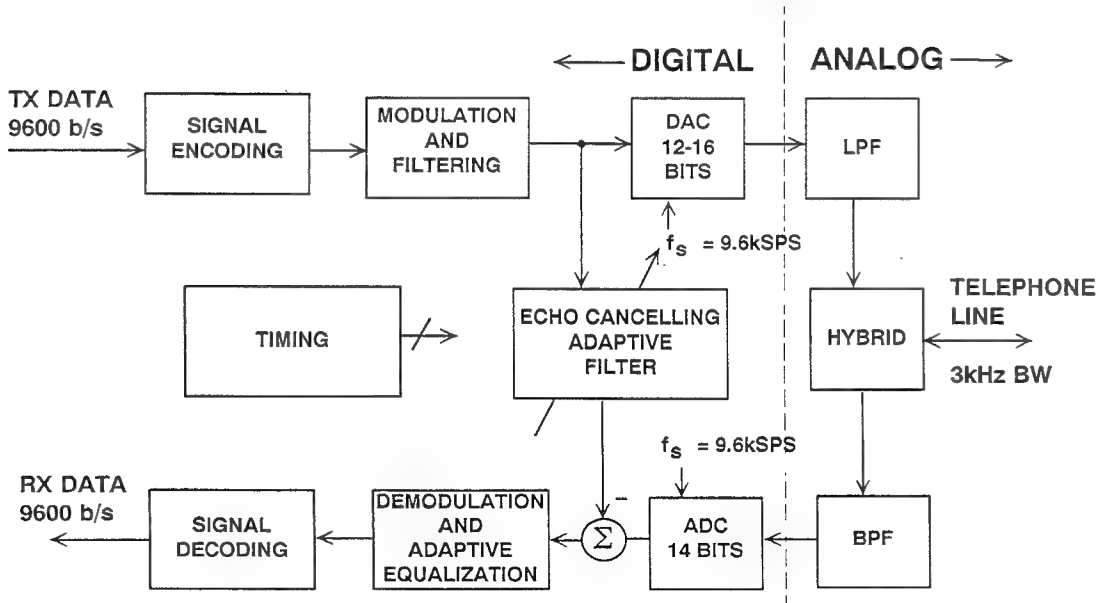


Figure 10.6

V.32 MODEM TRANSMITTER

A block diagram of the V.32 transmitter is shown in Figure 10.7. The input serial bit stream is first scrambled. Scrambling takes the input bit stream and produces a pseudo-random sequence. The purpose of the scrambler is to whiten the spectrum of the transmitted data. Without the scrambler, a long series of identical symbols could cause the receiver to lose carrier lock. Scrambling makes the transmitted spectrum resemble white noise, to utilize the bandwidth of the channel more efficiently, makes carrier recovery and timing synchronization easy, and makes adaptive equalization and echo cancellation possible.

The scrambled bit stream is divided into groups of four bits. The first two bits of each 4-bit group are first differentially encoded and then convolutionally encoded. This produces a 5-bit trellis-coded symbol in which the extra bit is a redundantly coded bit.

The 5-bit symbols are then mapped into the signal space using trellis-coding as defined in the V.32 recommendation. The signal space mapping produces two coordinates, one for the real part of the QAM modulator and one for the imaginary part. A diagram of the resulting V.32 signal constellation is shown in Figure 10.8.

Used prior to modulation, the digital pulse shaping filters attenuate frequencies above the Nyquist frequency that are generated in the signal mapping process. These filters are designed to have zero crossings at the appropriate frequencies to cancel intersymbol interference. The pulse shape filter is based on the impulse response of a raised cosine function as shown in Figure 10.9. The value T is equal to the reciprocal of the symbol rate (2400 symbols/second). For a sampling rate of 9600Hz and a symbol rate of 2400Hz, a 17-tap FIR filter can be used.

V.32 MODEM TRANSMITTER

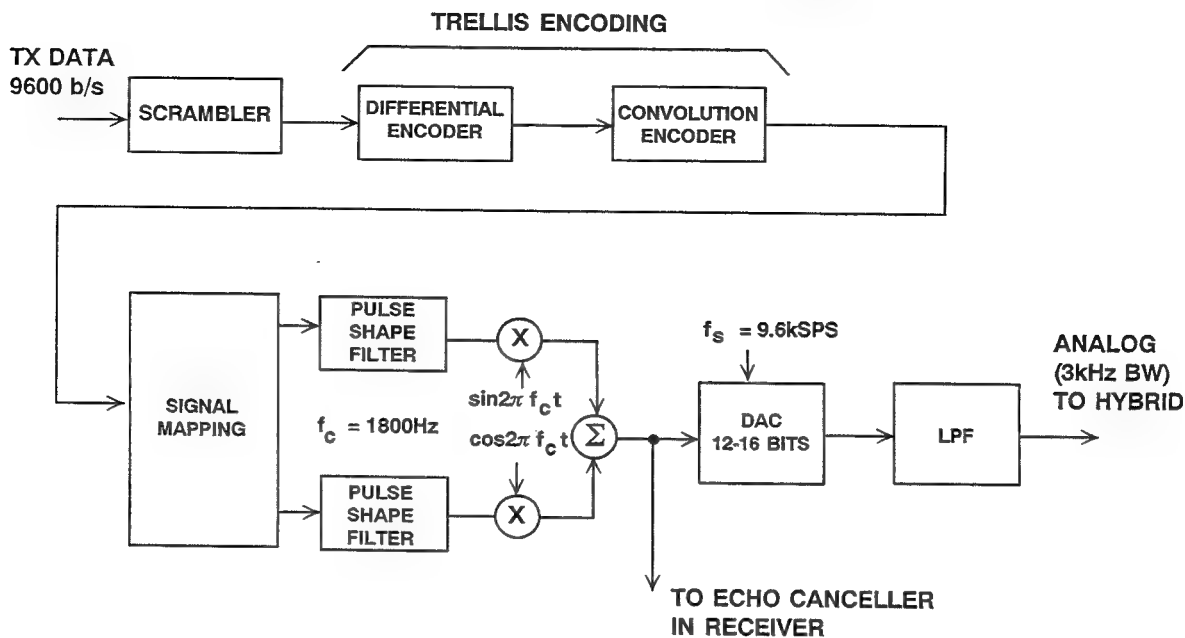


Figure 10.7

V.32 MODEM SIGNAL CONSTELLATION




Figure 10.8

The modulation for the V.32 coding scheme is quadrature amplitude modulation (QAM). Modulation is easily implemented in modern DSP processors. The process of modulation requires the access of a sine or

cosine value, the access of an input symbol (x or y coordinate) and a multiplication. The parallel architecture of the ADSP-2101 permits all three operations to be performed in a single 80ns cycle.

PULSE SHAPING FILTER IMPULSE RESPONSE



The diagram illustrates a digital receiver system. The input is an **ANALOG (3kHz BW)** signal **FROM HYBRID**. This signal passes through a **BPF** (Band Pass Filter) and then an **ADC 14 BITS**. The sampling rate is $f_s = 9.6\text{ kSPS}$. The ADC output is fed into a **Σ** (summing junction). The summing junction also receives a signal from the **ECHO CANCELLING ADAPTIVE FILTER** (labeled with a **-** sign) and a feedback signal from the **DESCRAMBLER**. The output of the summing junction is fed into the **ECHO CANCELLING ADAPTIVE FILTER** (labeled with a **+** sign). The output of the filter is fed into the **ADAPTIVE EQUALIZER**. The **ADAPTIVE EQUALIZER** also receives a feedback signal from the **DESCRAMBLER**. The output of the **ADAPTIVE EQUALIZER** is fed into the **VITERBI DECODER**. The **VITERBI DECODER** also receives a feedback signal from the **DESCRAMBLER**. The output of the **VITERBI DECODER** is fed into the **DIFFERENTIAL DECODER**. The output of the **DIFFERENTIAL DECODER** is fed into the **DESCRAMBLER**. The **DESCRAMBLER** outputs **RX DATA 9600 b/s**. The **DESCRAMBLER** also receives a feedback signal from the **VITERBI DECODER**.

Figure 10.10

V.32 MODEM RECEIVER

A block diagram of the V.32 modem receiver is shown in Figure 10.10. The receiver is made up of several functional blocks: the input antialiasing filter and ADC, a demodulator, an adaptive equalizer, a Viterbi decoder, an echo canceller, a differential decoder, and a descrambler. The receiver DSP algorithms are both memory-intensive and computation-intensive. The ADSP-2101 addresses both needs, providing 2K of program memory RAM (for both code and data) on chip, 1K of data memory RAM on chip and an instruction execution rate of 12.5MIPS.

The antialiasing filter and ADC in the receiver need to have a dynamic range from the largest echo signal to the smallest. The received signal can be as low as -40dBm, while the near-end echo can be as high as -6dBm. In order to insure that the analog front end of the receiver does not contribute any significant impairment to the channel under these conditions, an instantaneous dynamic range of 84dB (14 bits) and an SNR of 72dB is required.

In order to compensate for amplitude and phase distortion in the telephone channel, equalization is required to recover the transmitted data at an acceptably low bit error rate. In order to respond to rapidly changing conditions on the telephone line, adaptive equalization is required for the V.32 modem receiver. An adaptive equalizer can be implemented digitally in an FIR filter whose coefficients are continuously updated based on current line conditions. A 64-tap fractionally spaced equalizer provides the performance necessary for V.32 applications.

Separation between the transmit and receive signal in the V.32 modem is accomplished using echo cancellation. Echo cancellation is mandatory since both the calling and the answering modem use the same carrier frequency of 1800Hz. Both near-end and far-end echo must be cancelled in order to yield reliable communication. Echo cancellation is achieved by subtracting an estimate of the echo return signal from the actual received signal. The predicted echo is

AD7869 14-BIT, 83kSPS I/O PORT BLOCK DIAGRAM

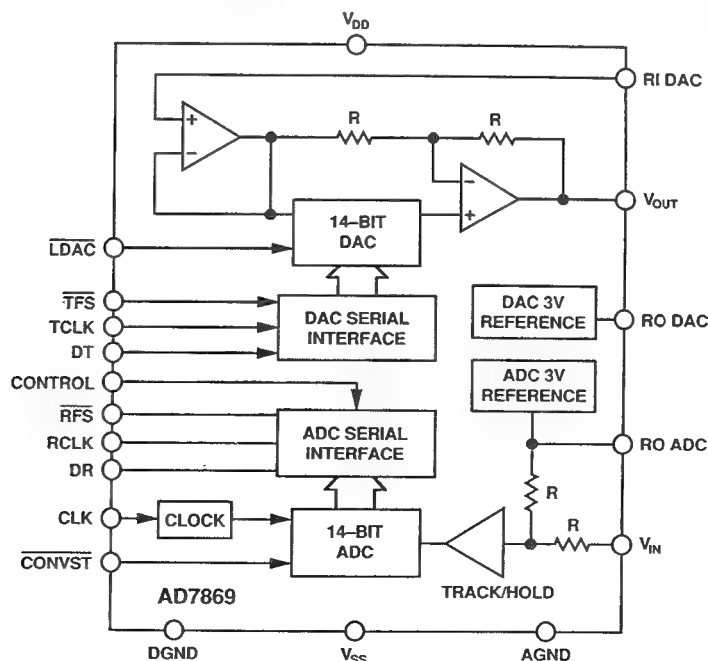


Figure 10.11

determined by feeding the transmitted signal into an adaptive filter with a transfer function that approximates the telephone channel. The adaptive filter commonly used in echo cancellers is the FIR filter (chosen for its stability and linear phase response), where the taps are determined using the least-mean-square (LMS) algorithm during a training sequence executed prior to full-duplex communications. The echo canceller must be able to cancel 16ms of echo. At 9600 samples/second, a 154-tap FIR filter is required to cancel the echo. Assuming that the canceller and frequency shifter have converged during the training period, about 200 cycles are required to cancel an echo in a V.32 modem.

The most common technique for decoding the received data is Viterbi decoding. Named after its inventor, the Viterbi algorithm is a general-purpose technique for making an

error-corrected decision. Viterbi decoding provides a certain degree of error correction by examining the received bit pattern over time to deduce the value that was the most likely to have been transmitted at a particular time. Viterbi decoding is computation-intensive. A history for each of the possible symbols sent at each symbol interval has to be maintained. For the V.32 modem, the symbol history spans 20 symbol intervals. At each symbol interval, the length of the path backward in time from each possible received symbol to a symbol sent some time ago is calculated. After 20 symbol intervals, the symbol that has the shortest path back to the original signal is chosen to be the current decoded symbol. A complete description of Viterbi decoding and its implementation on the ADSP-2100 family of DSP processors is given in Reference 2.

V.32 ANALOG FRONT END

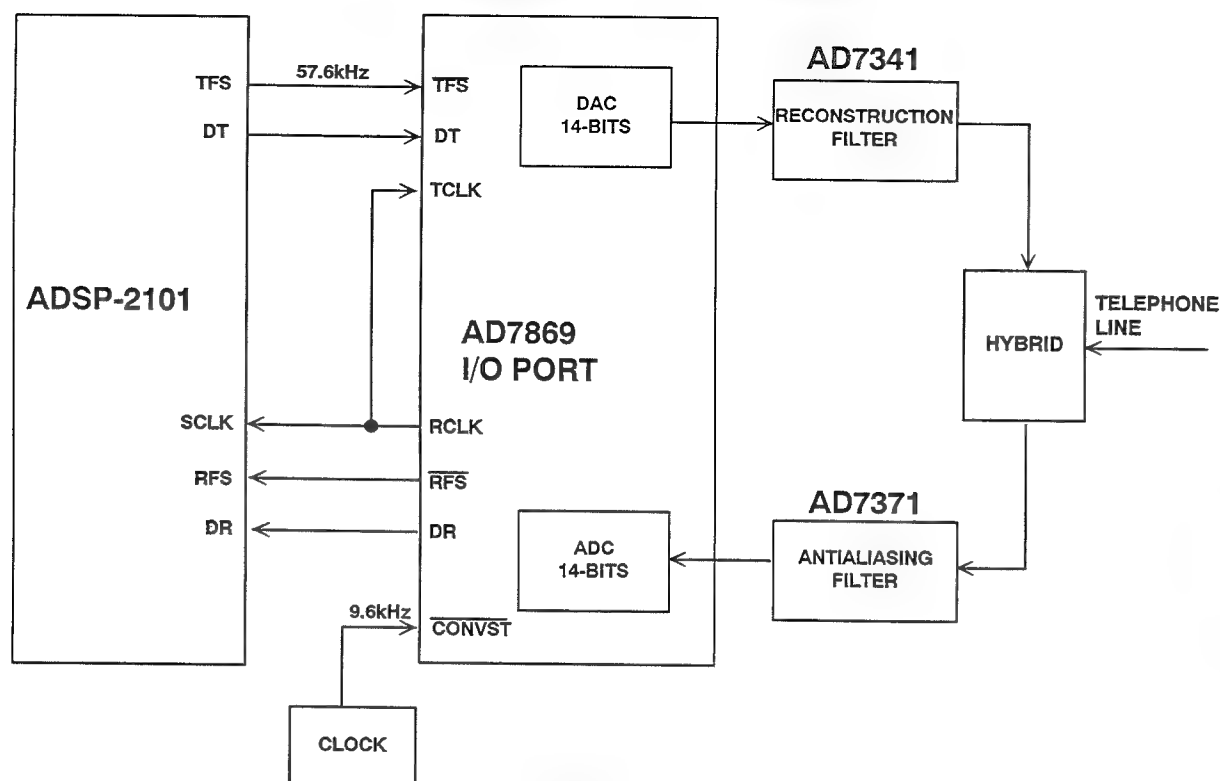


Figure 10.12

I/O PORTS AND CODECS FOR V.32 MODEMS

The AD7869 is a complete 14-bit, 83kSPS I/O Port with a zero-chip serial interface to most DSP processors such as the ADSP-2101, TMS3020/C25, and DSP56000. The SNR (including distortion) of the AD7869 is 80dB which meets the requirements for the V.32 transmit and receive channels. A block diagram of the device is shown in Figure 10.11

A block diagram of a complete analog front-end for a V.32 modem is shown in Figure 10.12. The serial interface is shown with the ADSP-2101 DSP processor.

The AD7341/AD7371 is a switched capacitor voiceband reconstruction/antialiasing filter chip set designed to be used in conjunction with the AD7869 to implement a complete V.32 modem analog front end. The

SCFs are clocked at a 57.6kHz rate, representing an oversampling ratio of 6X with respect to the ADC sampling rate of 9.6kSPS. The AD7341 is the transmit reconstruction filter. It implements the filter function using a seventh-order lowpass SCF and a second-order continuous time filter. The cutoff frequency is 3.5kHz. The AD7371 is the receive ADC antialiasing filter. It is a band-pass filter with a lower cutoff frequency of 180Hz and an upper cutoff frequency of 3.5kHz. The filter function is implemented using a second-order lowpass continuous time filter, a fourth-order highpass SCF and a seventh-order lowpass SCF. Key specifications for the AD7341/AD7371 SCF chip set are summarized in Figure 10.13.

KEY SPECIFICATIONS FOR THE AD7341/AD7371 MODEM FILTER CHIP SET

- **Stopband Attenuation:** 70dB, $f \geq 6.1\text{kHz}$
40dB, $f \leq 60\text{Hz}$
- **In-Band Signal to Noise Ratio:** 75dB
- **Total Harmonic Distortion:** -75dB
- **Differential Group Delay:** 350 μs
- **Programmable Attenuation (AD7341) :** 0 to 38dB
- **Programmable Gain (AD7371):** 0 to 24dB

Figure 10.13

The ADSP-28msp01 is a complete analog front end for high performance modems. The device has an architecture similar to the ADSP-28msp02 voiceband codec. The ADSP-28msp01 contains a 16 bit sigma-delta ADC and DAC and is capable of sampling rates of

7.2, 8.0, and 9.6kSPS with SNR and THD performance of 84dB. The extensive support of bit, baud, and convert clocks allow the ADSP-28msp01 to support many modem standards such as the V.32. Key specifications are summarized in Figure 10.14.

ADSP-28msp01 INTEGRATED MODEM ANALOG FRONT END KEY SPECIFICATIONS

- 16 Bit Sigma-Delta ADC and DAC
- On-Chip Antialiasing and Anti-Imaging Filters
- On-Chip Clock Generation Circuitry
- 84 dB THD and SNR
- Programmable Sampling Frequency of 7.2, 8.0, and 9.6kSPS
- DSP Compatible Serial Port
- 28 Pin DIP/SOIC

Figure 10.14

DIGITAL MOBILE RADIO

OVERVIEW

The rapidly growing number of cellular mobile phones in the United States has created significant system performance problems, especially in crowded metropolitan areas such as New York and Los Angeles. Call blocking during rush hour, flaws in call processing (disconnects and misconnects), and undesirable interchannel crosstalk are only a few. In addition, the current system lacks privacy and security, and data trans-

mission over a mobile link is almost impossible at rates above 1200 bits/s. These factors have led to the search for a more efficient and robust system based on digital techniques. Several digital approaches are being considered in the United States, while the Pan-European Digital Cellular Radio System (also known as *Groupe Speciale Mobile*, or GSM) has been defined and will be introduced in Europe in 1991.

PROBLEMS WITH CURRENT ANALOG CELLULAR RADIO

- Call Blocking During Busy Hours
- Misconnects and Disconnects due to Rapidly Fading Signals
- Lack of Privacy and Security
- Data Transmission Limited to 1200 bits/s

Figure 10.15

The current system in the United States is a cellular system based on Frequency Division Multiple Access (FDMA). A region is broken up into cells, with each cell having its own base station and its own group of assigned frequencies (see Figure 10.16). Be-

cause the radius of each cell is small (10 miles, for example) low power transmitters and receivers can be used. The cellular system lends itself to frequency reuse, since cells which are far enough apart can utilize the same band of frequencies without inter-

CELLULAR RADIO FREQUENCY REUSE




Figure 10.16

ference. The base stations must be linked together with an elaborate central control network so that a call may be handed-off to another cell when the signal strength from the mobile unit becomes too low for the current cell to handle.

The frequency spectrum allocation for cellular radio in the United States is approximately 825 to 850MHz and 870 to 895MHz. Conventional architectures (both analog and digital) are channelized. The total spectrum is divided up into a large number of relatively narrow channels, defined by a carrier

frequency. The carrier frequency is frequency-modulated with the voice signal using analog techniques. Each full-duplex channel requires a pair of frequencies, each with a bandwidth of approximately 30kHz. A user is assigned both frequencies for the duration of the call. The forward and reverse channel are widely separated, to help the radio keep the transmit and receive functions separated. The 40MHz allocated to cellular service can therefore be divided up into 666 frequency pairs, each serving one full-duplex circuit.

FREQUENCY DIVISION MULTIPLE ACCESS MOBILE RADIO SYSTEM

- Uses 825-850MHz and 870-895MHz Spectrum
- 30kHz Transmit, 30kHz Receive
- Analog Frequency Modulation (FM)
- Approximately 700 Users Capacity

Figure 10.17

Time Division Multiple Access (TDMA) allocates bandwidth on a time-slot basis. In the proposed United States TDMA system, the entire 30kHz channel is assigned to a particular transmission, but only for a short period of time. A 3:1 multiplexing scheme means that three conversations can take place with TDMA using the same amount of bandwidth as one analog cellular conversation does. Each transmit/receive sequence occurs on time slots lasting 6.7ms. The TDMA system relies on an extensive amount of DSP technology to reduce the coded speech bit-rate as well as to prepare the digital data for transmission over the analog medium. The TDMA approach has been chosen for the Pan-European GSM system and will be discussed later in more detail.

The second digital approach being considered in the United States is called Code Division Multiple Access (CDMA). This technique has been used in secure military

communications for a number of years under the name of *spread spectrum*. In spread spectrum, the transmitter transmits in a pseudo-random sequence of frequency hops over a relatively wide frequency range. The receiver has access to the same random sequence and can decode the transmission. The effect of adding additional users on the system is to decrease the overall signal to noise ratio for all the users. With this technique, the effect of allowing more calls than the normal capacity is to increase the bit-error rate for all users. New callers can keep coming in, interference levels will rise gradually, until at some point the process will become self-regulating: the quality of the voice link will become so bad that users will cut short or refrain from making additional calls. No one is ever blocked in the conventional sense, as they are in FDMA or TDMA systems when all channels or slots are full.

DIGITAL MOBILE RADIO APPROACHES

- Time Division Multiple Access (TDMA) - User Allocation Based on Time Slots: At Least 3X More Capacity than FDMA
- Code Division Multiple Access (CDMA) - Based on Spread Spectrum Technology: More Users Cause Graceful Degradation in Bit-Error Rate
- Both TDMA and CDMA Make Extensive Use of DSP in Speech Encoding and Channel Coding for Transmission

Figure 10.18

Both TDMA and CDMA systems make extensive use of DSP algorithms in both speech encoding and in preparing the signal for transmission. In the receiver, DSP techniques are used for demodulation and decoding the speech signal. The remainder of this

section will concentrate on speech processing and channel coding as they relate to the Pan-European GSM system. This will serve to illustrate the fundamental principles which are applicable to all digital mobile radio systems.

THE GSM SYSTEM

Figure 10.19 shows a simplified block diagram of the GSM Pan-European Digital Cellular Telephone System. The *speech encoder and decoder* and *discontinuous transmission* function will be described in detail. Up conversion and downconversion

portion of the system contain a digital modem similar to the V.32 recommendation previously discussed. Similar functions are performed digitally such as equalization, convolutional coding, Viterbi decoding, modulation and demodulation.

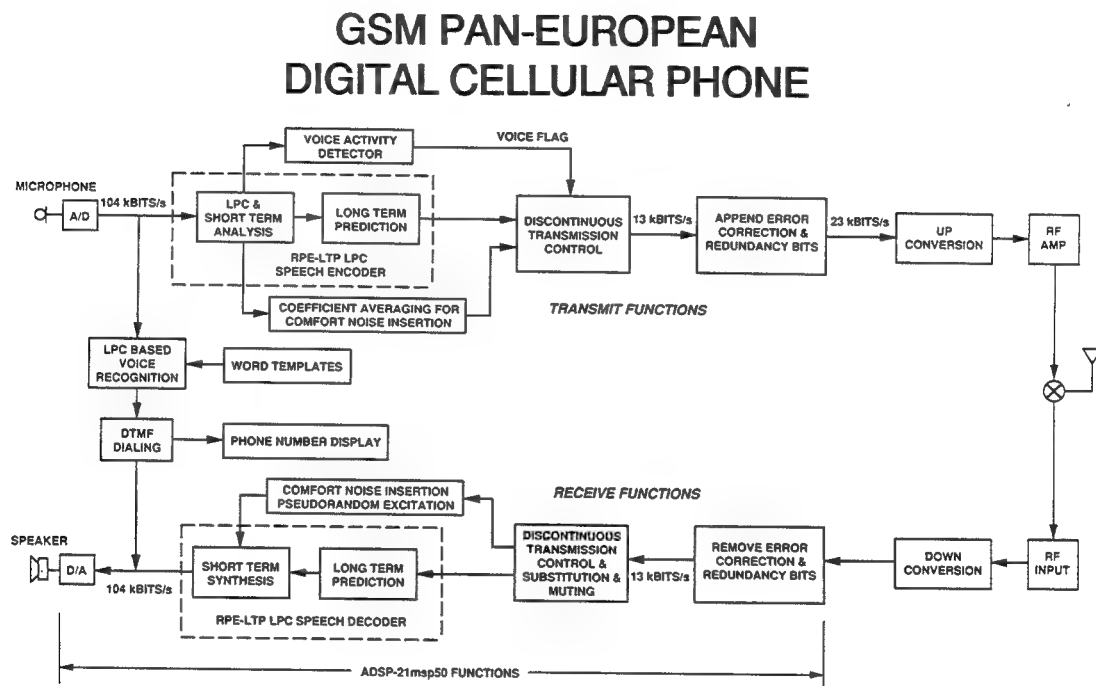


Figure 10.19

SPEECH CODEC

The standard for encoding voice has been set in the T-Carrier digital transmission system. In this system, speech is logarithmically encoded to 8 bits at a sampling rate of 8kSPS. The logarithmic encoding and decoding to 8 bits is equivalent to linear encoding and decoding to 13 bits of resolution. This produces a bit-rate of 104kb/s. The Speech Encoder portion of the GSM system compresses the speech signal to 13kb/s, and the

decoder expands the compressed signal at the receiver. The terms *codec* and *transcoder* are both often used to refer to the entire encoding and decoding speech compression function. The speech encoder is based on an enhanced version of linear predictive coding (LPC). The LPC algorithm uses a model of the human vocal tract that represents the throat as a series of concentric cylinders of various diameters. An excitation (breath) is

forced into the cylinders. This model can be mathematically represented by a series of simultaneous equations which describe the cylinders.

The excitation signal is passed through the cylinders, producing an output signal. In the human body, the excitation signal is air moving over the vocal cords or through a constriction in the vocal tract. In a digital system, the excitation signal is a series of pulses for vocal excitation, or noise for a constriction. The signal is input to a digital lattice filter. Each filter coefficient represents the size of a cylinder.

An LPC system is characterized by the number of cylinders it uses in the model. Eight cylinders are used in the GSM system, and eight reflection coefficients must be generated.

Early LPC systems worked well enough to understand the encoded speech, but often the quality was too poor to recognize the voice of the speaker. The GSM LPC system employs two advanced techniques that improve the quality of the encoded speech. These techniques are *regular pulse excitation* (RPE) and *long term prediction* (LTP). When these techniques are used, the resulting quality of encoded speech is nearly equal to that of logarithmic pulse code modulation (companded PCM as in the T-Carrier system).

The actual input to the speech encoder is a series of 13-bit samples of uniform PCM speech data. The sampling rate is 8kHz. The speech encoder operates on a 20ms window (160 samples) and reduces it to 76 coefficients (260 bits total), resulting in an encoded data rate of 13kb/s.

SPEECH COMPRESSION IN THE GSM SYSTEM

- Input Data: 13bit Samples at 8KSPS = *104kbits/s*
- Output Data for Each 20ms Window: 76 Filter Coefficients, 260 bits Total = *13kbits/s*

Figure 10.20

DISCONTINUOUS TRANSMISSION (DTX)

Discontinuous transmission (DTX) allows the system to shut off transmission during the pauses between words. This reduces transmitter power consumption and increases the overall GSM system's capacity.

Low power consumption prolongs battery life in the mobile station and is an important consideration for hand-held portable phones. Call capacity is increased by reducing the

interference between channels, leading to better spectral efficiency. In a typical conversation each speaker talks for less than 40% of the time, and it has been estimated that DTX can approximately double the call capacity of the radio system.

The required DTX functions are summarized in Figure 10.21.

DISCONTINUOUS TRANSMISSION (DTX) FUNCTIONS

- **Voice Activity Detection (VAD) to Detect Speech**
- **Comfort Noise Insertion (CNI) to Synthesize Artificial Car Noise During Pauses Between Words**
- **Output Muting When Lost Speech Frames Are Received**

Figure 10.21

The voice activity detector (VAD) is located at the transmitter; its job is to distinguish between speech superimposed on the background noise and noise with no speech present. The input to the voice activity detector is a set of parameters computed by the speech encoder. The VAD uses this information to decide whether or not each 20ms frame of the encoder contains speech.

Comfort noise insertion (CNI) is performed at the receiver. The comfort noise is generated when the DTX has switched off the transmitter; it is similar in amplitude and spectrum to the background noise at the transmitter. The purpose of the CNI is to eliminate the unpleasant effect of switching between speech with noise, and silence. If you were listening to a transmission without CNI, you would hear rapid alternating between speech in a high-noise background (i.e. in a car), and silence. This effect greatly reduces the intelligibility of the conversation.

When DTX is in operation, each burst of speech is transmitted followed by a *silence descriptor* (SID) frame before the transmission is switched off. The SID serves as an end of speech marker for the receive side. It contains characteristic parameters of the background noise at the transmitter, such as spectrum information derived through the use of linear predictive coding.

The SID frame is used by the receiver's comfort noise generator to obtain a digital

filter which, when excited by pseudo-random noise, will produce noise similar to the background noise at the transmitter. This comfort noise is inserted into the gaps between received speech bursts. The comfort noise characteristics are updated at regular intervals by the transmission of SID frames during speech pauses.

Redundant bits are then added by the processor for error detection and correction at the receiver, increasing the final encoded bit rate to 22.8kb/s. The bits within one window, and their redundant bits, are interleaved and spread across several windows for robustness.

The ADSP-21msp50 Mixed Signal Processor shown in Figure 10.22 can perform all of the above tasks within the 20ms sampling window because of its optimized DSP architecture and the special on-chip peripherals associated with it. The sigma-delta converters provide the necessary interface to the speaker and microphone. The parallel host interface port communicates with a host processor, which is responsible for loading the ADSP-21msp50 with the appropriate programs during power-up, dialing, and actual conversation phases of a complete call. The ADSP-21msp50 has 1K words of (16-bit) data memory static RAM and 2K words of (24-bit) program memory static RAM on chip. The device operates at a 13MHz clock rate and has a low power mode and a power down

ADSP-21msp50 BLOCK DIAGRAM

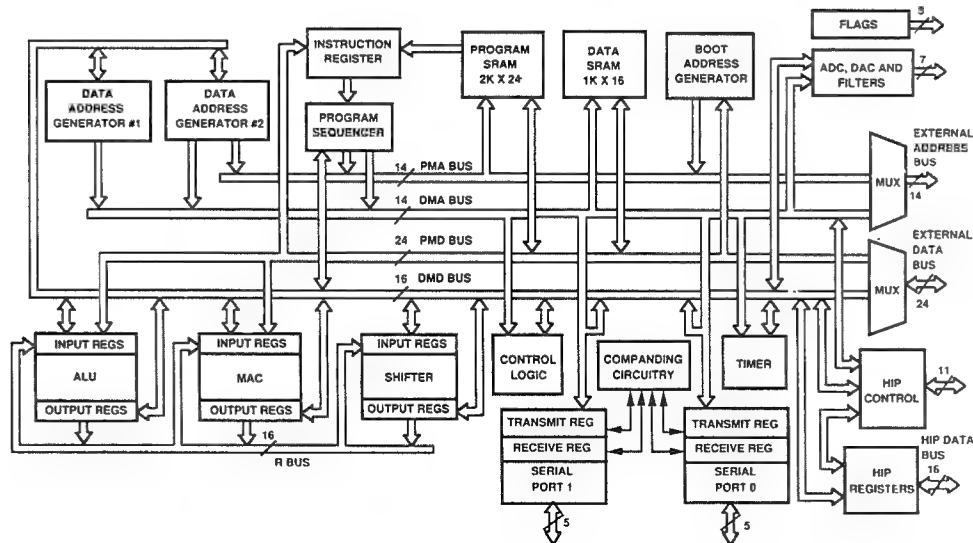


Figure 10.22

mode (less than 1mW in power down). The ADSP-21msp50 combines the core ADSP-2100 architecture (three computational units, data address generators, and a program sequencer) with two serial ports, a programmable timer, host interface port, an

analog interface, extensive interrupt capabilities. Key features of the ADSP-21msp50 are summarized in Figure 10.23, and the benchmark performance in the GSM system is shown in Figure 10.24.

ADSP-21msp50 MIXED SIGNAL PROCESSOR KEY SPECS

- On-Chip 16-Bit Sigma-Delta ADC and DAC
- 65dB SNR and THD
- 8kSPS Sampling Frequency, 1MHz Clock (125X Oversampling)
- 2K Words Program Memory Ram (24-bits)
- 1K Words Data Memory Ram (16-bits)
- 13MIPS Performance
- Host Interface Port
- ADSP-2100 Family Compatible Instruction Set
- Low Power and Power Down Mode

Figure 10.23

ADSP-21msp50 GSM BENCHMARKS

FUNCTION	CYCLE COUNT MAXIMUM WORST CASE	TIME REQUIRED OUT OF 20ms WINDOW	PROCESSOR LOADING
RPE-LTP LPC Encoder	49300	3.8ms	19.0%
RPE-LTP LPC Decoder	14400	1.1ms	5.5%
Voice Activity Detector	2141	0.17ms	0.9%
Total Functions	65841	5.07ms	25.4%
Free			74.6%

Internal Program Memory Required: 1988 words
Internal Data Memory Required: 964 words

Figure 10.24

GSM SYSTEM UPCONVERSION AND DOWNCONVERSION

A block diagram of the GSM system with particular emphasis on the upconversion and downconversion circuitry is shown in Figure 10.25. The transmit data coming from the speech processor contains error correction and redundancy bits. The bit rate at this point in the system is 23kb/s. The channel coder and filters prepare the data to fit the TDMA format of the GSM system. Figure

10.26 shows how each 200kHz of frequency spectrum contains data from 8 users. Each user is assigned a time slot of 0.577ms during which time a burst of 156 data bits are transmitted at a modulation frequency of approximately 270kHz. Modulation is accomplished using Gaussian Minimum Shift Keying (GMSK), a form of frequency shift keying which minimizes spectral leakage.

GSM BLOCK DIAGRAM: UPCONVERSION DETAILS

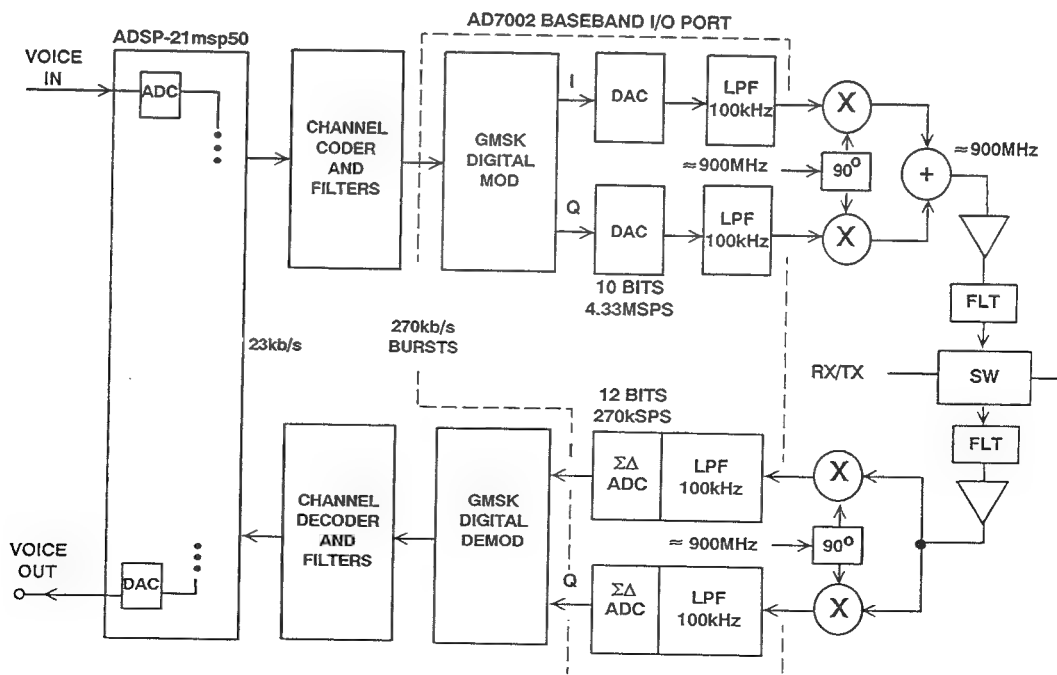


Figure 10.25

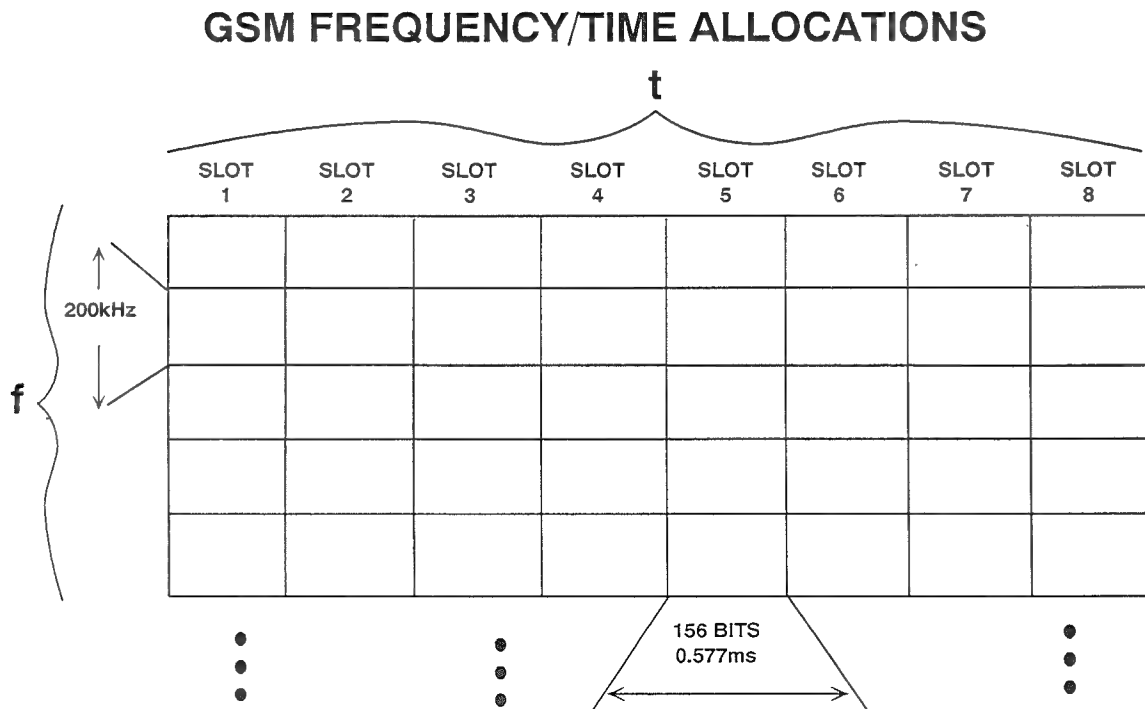


Figure 10.26

The modulation is done digitally and converted into an I and Q signal. The modulator outputs drive two 10-bit DACs whose filtered output drives the RF modulators. The DACs are oversampled by a factor of 16 in order to simplify the anti-imaging analog filter requirements. The combined I and Q signal drives the RF amplifier, filter, and the antenna.

The received signal is filtered, amplified, and fed to an I/Q RF demodulator which recovers the I and Q signals. The baseband I and Q signals are converted by two 12-bit DAC at an effective sampling rate of 270kSPS. The I and Q signals are then demodulated by the GMSK digital demodulator. The 270kb/s burst is sent to the channel decoder and filters and then to the speech processor.

The AD7002 is a complete GSM Baseband I/O Port which performs the functions shown in Figure 10.25. The transmit path contains two 10-bit oversampled (16X) DACs followed by fourth-order anti-imaging filters. The DACs are driven by a digital modulator containing a GMSK-coded ROM. The receive path contains two high-performance 12 bit sigma-delta ADCs having a throughput rate of 270kSPS. The sigma-delta ADCs contain a 288-tap FIR filter having linear phase response and a 3dB point of 122kHz. Three auxiliary DACs are included for such functions as AFC, AGC and carrier shaping. The device dissipates approximately 100mW and has flexible power-down or sleep modes. Key specifications are summarized in Figure 10.27.

AD7002 GSM BASEBAND I/O PORT KEY SPECIFICATIONS

- **Transmit Path:** **GMSK I/Q Digital Modulator**
 Dual 10 Bit, 4.33MSPS Oversampled DACs
 Dual Anti-Imaging Filters
- **Receive Path:** **Dual 12 Bit 270kSPS Sigma-Delta ADCs**
 288-Tap 100kHz Linear Phase FIR Filter
- **3 Auxiliary DACs for AFC, AGC**
- **Low Power: 100mW**
- **Sleep Mode**

Figure 10.27

DIGITAL AUDIO STUDIO RECORDING

The activities related to studio recording are complex and varied. Generally, multiple channels are used, with each track dedicated to one or more sources (instruments/voices). All channels need not be recorded at the same time. Each channel is subjected to extensive processing such as gain control, filtering, non-linear compression or expansion, reverberation, spectral equalization, and other special-effects enhancements. The contributing channels are then mixed together to obtain a final arrangement with the desired overall effect.

Traditionally, channel processing and mixing were implemented entirely in the analog domain—with numerous disadvan-

tages. Each channel's information—stored as an analog signal on magnetic tape—degrades as the cutting, splicing, and re-recording process progresses, undermining the benefits of the processing. The limited performance range available with analog processing sets a ceiling on the signal enhancement that can be obtained. Also, analog circuitry can only handle one channel at a time; multi-channel mixers are expensive and difficult to control. Finally, if analog processing hardware is used, overall mixing flexibility can be achieved only through hardware modifications. In practice, this means that the mixing process loses its ability to creatively explore special effects.

DIGITAL AUDIO STUDIO TECHNIQUES

- Digital Recording: 16, 18 or 20 Bits for ADC
- Digital Mixing
- Gain Control
- Reverberation and Special Effects
- Equalization using Digital Filters

Figure 10.28

Increasingly, audio processing is relying on digital techniques to improve audio quality. The first step in this transition was digital recording, which became prevalent in the early 1980s. Audio signals are first converted to digital form before being stored on magnetic tape. Digital recording eliminates several sources of degradation that hamper analog recordings, including the effects of non-linearities and additive noise in the magnetic materials used for recording, and wow and flutter in the tape playback mechanism.

In studio mixing applications, however, digital recording does not eliminate all complications. In the mixing and enhancement process, information is passed from one tape

to another—requiring both ADCs and DACs, a source of noise. These conversions are no longer necessary if all processing and mixing are handled with DSP techniques.

In the DSP-based studio recording system shown in Figure 10.29, signals are converted to digital as early as possible, usually to 16, 18, or 20-bit resolution. After conversion, the audio processing is handled digitally with high performance DSP processors. Gain factors are handled with digital multiplication. Filtering and equalization can be handled with linear-phase FIR filters. Dynamic-range control is easily included in the system by using a multiplier for non-linear compression/expansion computations.

DIGITAL AUDIO STUDIO SYSTEM

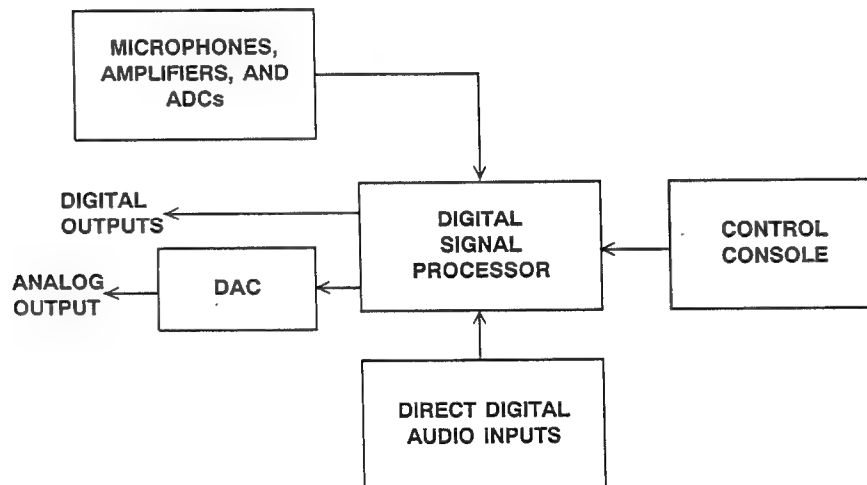


Figure 10.29

The traditional mixing process is also easily implemented in a DSP-based system. Digital channels to be mixed are simply added together. Relative time delay lags can be easily introduced into the channel flows, allowing phase delays to be equalized. Channel interconnections—which have to be hardwired in an analog processor—can be easily reconfigured in a DSP system.

In addition to improving on traditional operation, a DSP studio recording system opens up numerous new options. Unusual special effects are readily included in the system. Reverberation effects can be modeled, simulated, and integrated into the final recording. Digital reverberation can give concert hall or cathedral ambience to what might have been recorded in a dry studio. An FFT routine's spectral analysis of the signal forms the basis

for adaptive digital filters that provide optimal equalization.

With the advent of compact disc (CD) and digital audio tape (DAT) players, there is no requirement for digital-to-analog conversion anywhere in the studio recording process, except for monitoring and optimization purposes. The final digital recording can be transferred directly to the CD or DAT in digital form with no loss in fidelity.

Although 18 and 20-bit ADCs may be used in the recording process, the standard for CD and DAT has been set at 16 bits. Additional bits may be used in the DSP studio processing to allow for roundoff errors, overflows, etc., but the final recording is truncated to 16 bits per sample on the CD or DAT. The sampling-rate standard for CD recordings is 44.1kSPS, and 48kSPS for DAT.

DIGITAL AUDIO RECORDING STANDARDS

- **16-20 bits ADC Resolution, Truncated to 16 bits for Compact Disc**
- **44.1kSPS Sampling Rate for CD Players**
- **48kSPS Sampling Rate for Digital Audio Tape (DAT) Players**

Figure 10.30

Performance of audio systems is primarily measured in terms of three dynamic specifications: Total harmonic distortion plus noise

(THD+N), D-Range distortion, and signal-to-noise ratio. The definitions of these specifications are given in Figure 10.31.

KEY AUDIO PERFORMANCE SPECIFICATIONS

- **THD + N:** Ratio of the Square Root of the Sum of the Squares of the Values of the Harmonics and Noise to the Value of the Fundamental Input Frequency Expressed in % or dB
- **D-Range Distortion:** Ratio of the Distortion Plus Noise to the Signal at a Signal Amplitude of -60dB. Add 60dB to the Ratio to Obtain D-Range Distortion Value
- **Signal-to-Noise Ratio:** Ratio of the Amplitude of the Output with No Signal Present to the Amplitude of the Output When a Fullscale Output is Present

Figure 10.31

Specifications for the AD1876 16-bit 100kSPS ADC are given in Figure 10.32, and specifications for the AD1879 18-bit sigma-

delta ADC are given in Figure 10.33. These two ADCs are fully specified in terms of digital audio parameters.

AD1876 16-BIT 100kSPS SAMPLING ADC

- Autocalibrating
- 0.001% THD (100dB)
- 90dB S/(N + D)
- 2x Audio Oversampling Capability
- Power Dissipation: 250mW

Figure 10.32

AD1879 DUAL 18-BIT SIGMA-DELTA AUDIO ADC

- Signal-to-Noise Ratio: -106dB FS
- 0.0017% THD (95dB) at 1kHz
- Interchannel Crosstalk: -110dB at 1kHz
- Decimator Filter Passband Ripple: 0.001dB
- Decimator Filter Stopband Attenuation: 115dB
- Oversampling Ratio: 64x
- Output Word Rate: 48kSPS

Figure 10.33

COMPACT DISC (CD) PLAYER ELECTRONICS

A simplified block diagram of the read electronics for a typical CD player is shown in Figure 10.34. The read electronics takes the data from the CD read head and performs the necessary data qualification, error detection and error correction. Data from the read electronics is in serial format, 16 bits per sample, at an effective sampling rate of 44.1kHz per channel. Data for the two channels is usually multiplexed in a single

1.4112MHz bit stream. In theory, it is possible to reconstruct the audio signal using two 16-bit DACs preceded by a digital demultiplexer and parallel-to-serial converters operating at an update rate of 44.1kHz followed by analog anti-imaging filters. First-generation CD players used this approach as shown in Figure 10.35. An alternative is a single 16 bit DAC with output multiplexing for left and right channel.

COMPACT DISC PLAYER READ ELECTRONICS

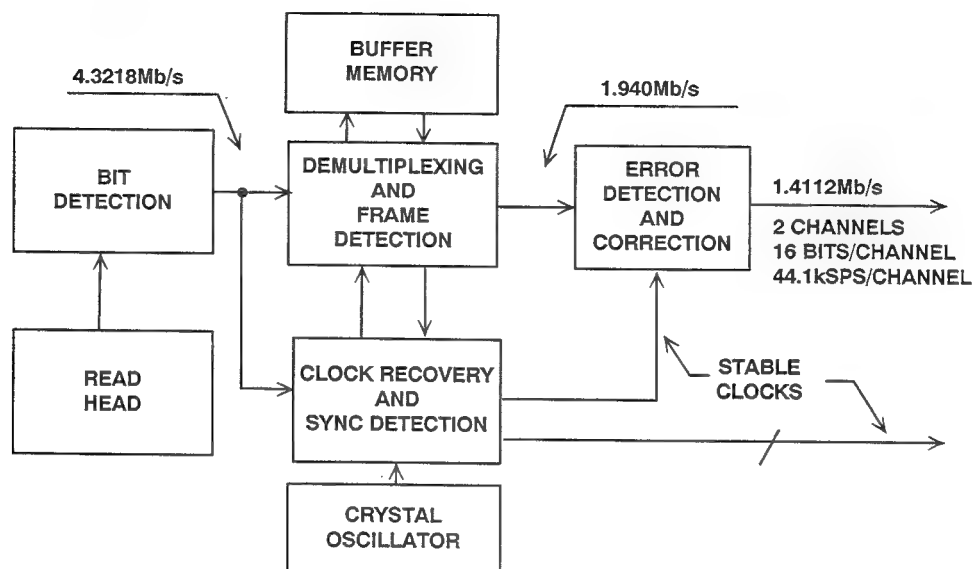


Figure 10.34

FIRST-GENERATION CD RECONSTRUCTION ELECTRONICS

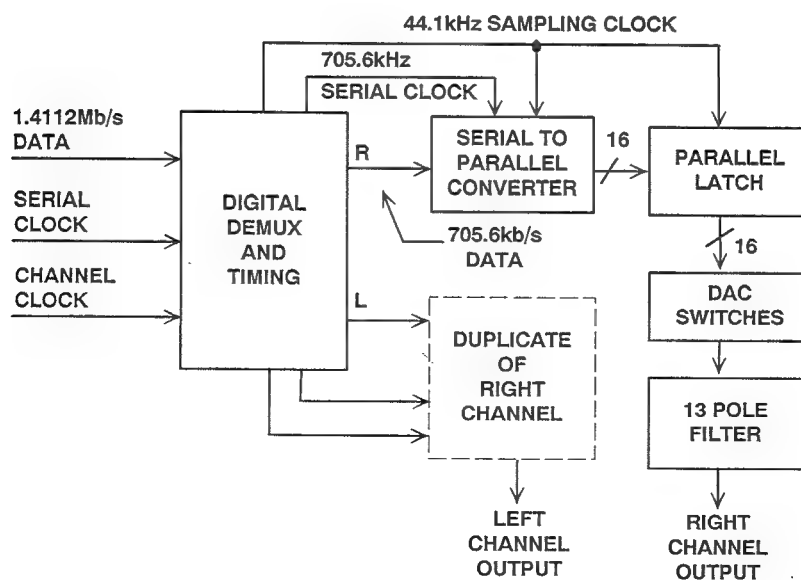


Figure 10.35

8X OVERSAMPLED 18-BIT CD RECONSTRUCTION ELECTRONICS

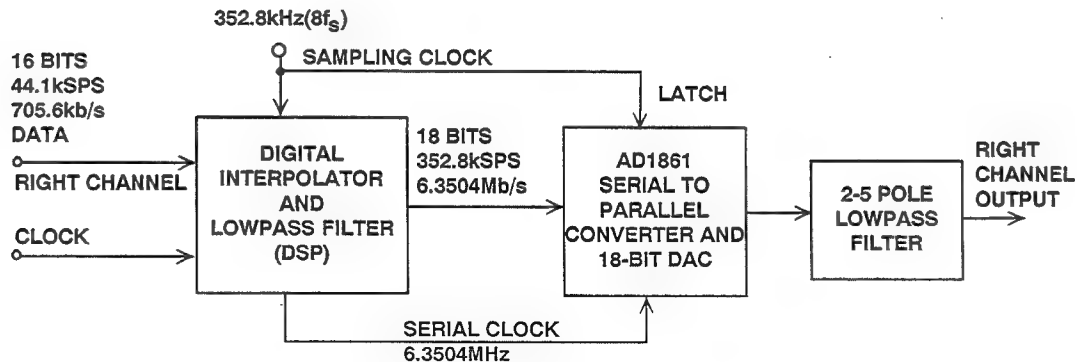


Figure 10.36

Sampling at 44.1kHz places severe requirements on the antialiasing filters. The audio bandwidth extends from 20Hz to 20,000Hz, and the filters must exhibit a flat frequency response over this frequency. In order to prevent aliasing, the filters must have at least 40dB to 50dB attenuation at 22.05kHz. This implies a complicated and costly 9- to 13-pole analog filter. In addition, these higher-order filters typically have non-linear phase response which is undesirable in audio applications. For this reason, the principles of oversampling and digital filtering are now in widespread use to simplify the design of the analog filter as well as increase the overall signal-to-noise ratio.

Second-generation CD players typically used oversampling ratios of 2x (88.2kSPS) or 4x (176.4kSPS) in conjunction with linear phase FIR digital interpolation filter chips. Third-generation players are using 8x oversampling (352.8kSPS) as shown in Figure 10.36, and the trend for future players will probably be 16x (705.6kSPS) or higher.

In addition to easing the requirements on the anti-imaging filter, oversampling followed by digital filtering spreads the quantization noise over a wider bandwidth, giving an improvement in SNR of $10 \log_{10}(K)$, where K is the oversampling ratio. This implies that for an oversampling ratio of 8x, there is a theoretical 9dB (or 1.5bits) improvement in SNR. It is possible to carry the arithmetic in the digital interpolation filter out to 18 bits, drive an 18-bit audio DAC with the result, and realize this improvement in practical CD players. If 16x oversampling were used, the theoretical improvement in SNR would be 12dB, or 2 bits (See Figure 10.37). A block diagram of a the complete reconstruction channel of an 8x oversampled 18 bit CD player is shown in Figure 10.38. The design is based on the dual 18 bit AD1865 DAC. Because of the 8x oversampling ratio, the output filter is the simple 3-pole filter which is shown in Figure 10.39.

EFFECTS OF OVERSAMPLING AND DIGITAL FILTERING ON CD PLAYER DESIGN

OVERSAMPLING RATIO K	THEORETICAL INCREASE IN SNR	USEFUL BITS OF DAC RESOLUTION	NUMBER OF POLES REQUIRED IN ANALOG FILTER
1	0dB	16	10
2	3dB	16	5
4	6dB	16/18	4
8	9dB	16/18/20	3
16	12dB	16/18/20	2

Figure 10.37

8X OVERSAMPLED CD RECONSTRUCTION ELECTRONICS USING DUAL 18-BIT DAC

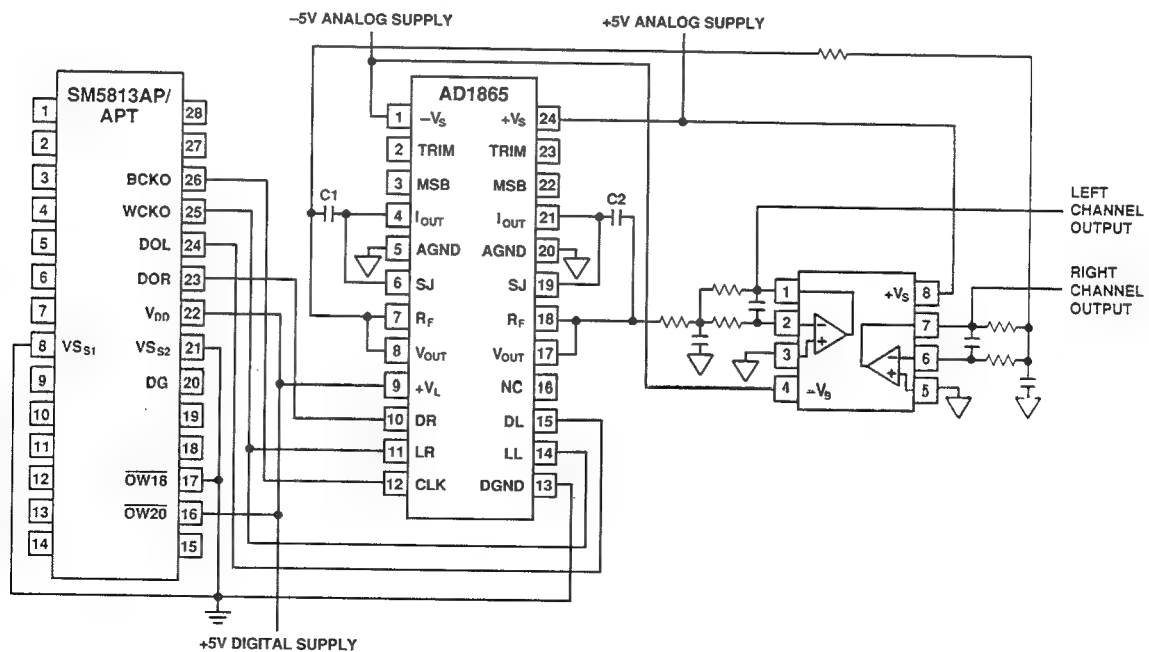


Figure 10.38

3-POLE ANTIALIASING FILTER FOR 18-BIT, 8X OVERSAMPLING

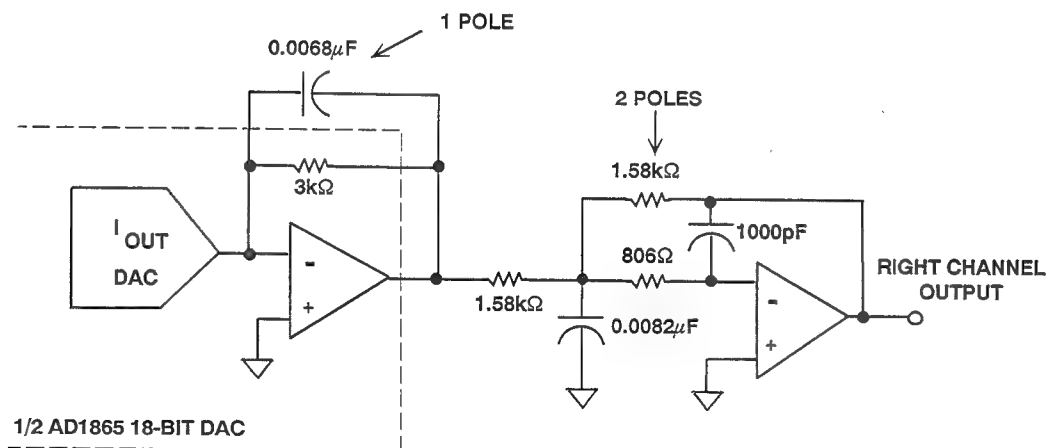


Figure 10.39

HIGH PERFORMANCE 20-BIT 8X OVERSAMPLING CD RECONSTRUCTION ELECTRONICS

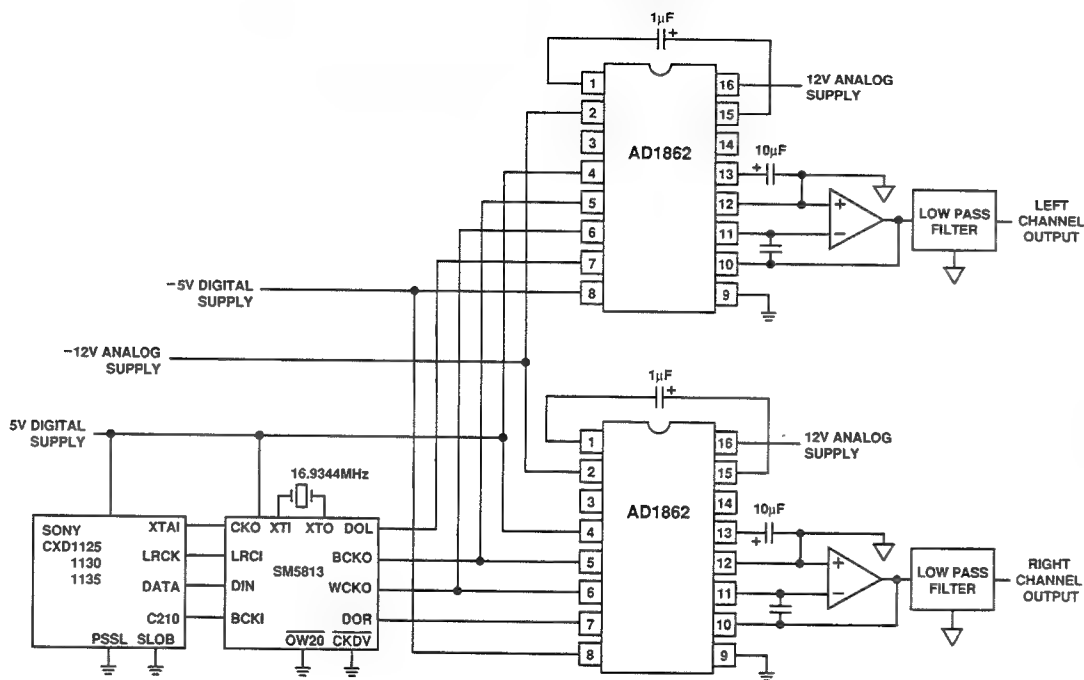


Figure 10.40

An additional reason for using DACs with greater than 16-bit resolution is that the process of digital interpolation and filtering adds truncation noise when the digital filter rounds off the interpolated values. This noise is reduced by using 18- and even 20-bit DACs to preserve accuracy in the interpolated values. A block diagram of a 20-bit, 8x oversampling CD filter and DAC configuration is shown in Figure 10.40. Because of the 8x oversampling ratio, a 5-pole lowpass filter is sufficient to maintain the required performance. Typical THD + N performance for the system shown in Figure 10.40 (including the output filter) is shown in Figure 10.41. A variety of CD digital interpolation filter chips are currently available from manufacturers such as Yamaha, NPC, and Sony.

There are a number of audio DACs currently available on the market ranging from

16 to 20-bit resolution. Newer devices are capable of sampling rates up to 768kSPS, allowing 16x oversampling. Unlike traditional DACs, audio DACs are specified in terms of ac parameters such as THD + N, SNR, and D-Range Distortion because traditional dc specifications are not as critical for audio applications. All audio DACs accept serial inputs and have internal serial-to-parallel converters followed by a parallel latch. Two clock inputs are therefore required to operate an audio DAC. A serial clock is needed to strobe the serial data into the serial-to-parallel converter, and a latch-enable clock is required to strobe the parallel latch. A simplified block diagram of a typical digital audio DAC is shown in Figure 10.42.

THD+N PERFORMANCE OF 20-BIT 8X CD ELECTRONICS USING AD1862




Figure 10.41

TYPICAL DC AUDIO DAC (SINGLE CHANNEL, 18 BITS, 8X)

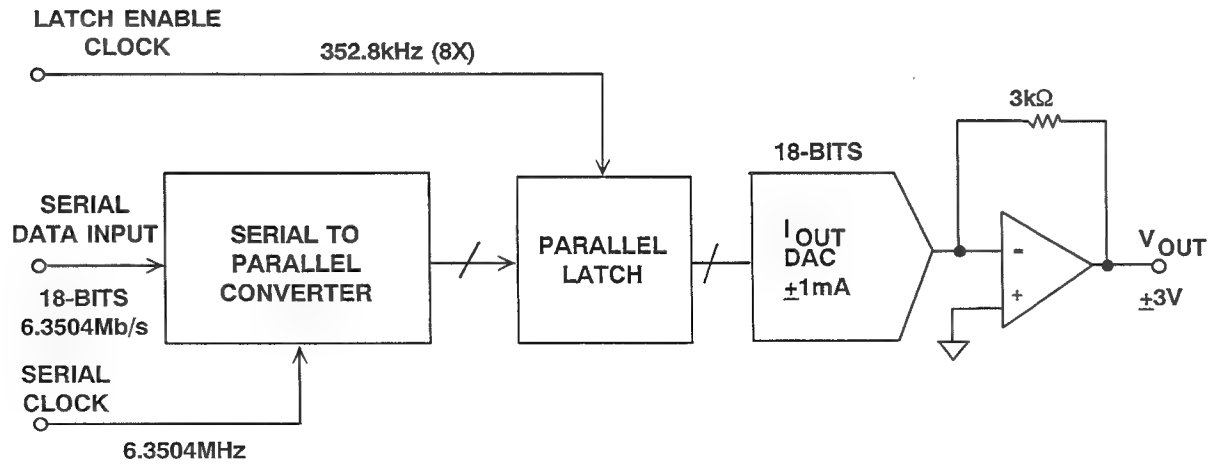


Figure 10.42

REFERENCES

1. John Bingham, **The Theory and Practice of Modem Design**, John-Wiley, 1988.
2. **ADSP-2100 Family Applications Handbook Vol. 4**, Analog Devices, 1990.
3. John Reidy and Mike Curtin, **Understanding and Applying the AD7341/AD7371 Switched Capacitor Filters**, Analog Devices Application Note E1373-15-5/90.
4. George Calhoun, **Digital Cellular Radio**, Artech House, Norwood MA, 1988
5. Nick Morley, *New Cellular Scheme Muscles In*, **EDN News**, November 15, 1990, p. 1.
6. P. Vary, K. Hellwig, R. Hofmann, R. J. Sluyter, C. Galand, M. Rosso, *Speech Codec for the European Mobile Radio System*, **ICASSP 1988 Proceedings**, p. 227.
7. D. K. Freeman, G. Cosier, C. B. Southcott, I. Boyd, *The Voice Activity Detector for the Pan-European Digital Cellular System*, **ICASSP 1989 Proceedings**, p. 369.
8. J. Haspeslagh, D. Sallaerts, et. al, *A 270-kb/s 35mW Modulator IC for GSM Cellular Radio Hand-Held Terminals*, **IEEE Journal of Solid State Circuits**, Vol. 25, No. 6, December, 1990, pp. 1450-1457.
9. Ken Pohlmann, *The Compact Disc Formats: Technology and Applications*, **J. Audio Eng. Soc.**, Vol. 36, No.4, April 1988, pp 250-280.
10. Gary Davis and Ralph Jones, **Sound Reinforcement Handbook**, Second Edition, Written for Yamaha, Hal Leonard Publishing Company, Milwaukee, 1989.

SECTION XI

MIXED SIGNAL CIRCUIT TECHNIQUES

MIXED SIGNAL CIRCUIT TECHNIQUES

■ INTRODUCTION

■ RESISTANCE

Resistance of Conductors

Skin Effect

Voltage Drop in Signal Leads-Kelvin Feedback

Leakage in Insulators: Guard Rings, Electrostatic Damage

**Parasitic Effects in Resistors: Inductance, Thermoelectric Effects
Stability and Matching, Voltage Variation of Resistance, Johnson Noise**

■ CAPACITANCE

Stray Capacitance

Faraday Shields

Noise

**Parasitic Effects in Capacitors: Capacitor Leakage, Series/Loss
Resistance, Inductance of Capacitors, Dielectric Absorption**

■ INDUCTANCE

Stray Inductance

Mutual Inductance

Ringling

Parasitic Effects in Inductors

Quality Factor (Q)

■ GROUNDING AND SIGNAL ROUTING

Signal Return Currents

Ground Noise and Ground Loops

Star (Mecca) Grounds

Separate Analog and Digital Grounds

Ground Plane

Transmission Lines

System Grounds

Signal Routing

■ POWER SUPPLIES

Power Supply Noise

Switching-Mode Power Supplies

■ ELECTROMAGNETIC INTERFERENCE

Radio Frequency Interference

Photoelectric Effects

■ LOGIC

Fan-Out

Timing Variations

Sampling Clock Noise

Logic Noise

■ PROBLEM AREAS

Limitations of SPICE Modelling

Sockets

Prototyping High Performance Analog Circuitry

SECTION XI

MIXED SIGNAL CIRCUIT TECHNIQUES

INTRODUCTION

There are considerably more problems involved in the successful design of Mixed Signal circuitry than mere circuit design. If we design an electronic circuit as a diagram, whether we use an old-fashioned pencil and paper or, as is the modern fashion, a computer and SPICE or some similar software, we are overlooking one of the most important factors in the design of successful hardware, namely that what we are designing is **HARDWARE**, and until it has been shown to work successfully in fact, rather than in simulation, our design is not complete.

This section of our seminar considers the problems which arise when reality reacts on a design which theory and modelling have shown to be satisfactory. What, in fact, has happened is that our model probably does not consider the effects of non-ideal components and of spurious or parasitic components resulting from the circuit layout which has been used. It is not, perhaps, too fanciful to describe this as the section of the Seminar dealing with **Murphy's Law**.

MURPHY'S LAW

IN ANY SET OF CIRCUMSTANCES THE WORST THING THAT CAN HAPPEN - WILL

- Any effect which you think can be disregarded, can't.
- Nature always sides with the hidden flaw.

Figure 11.1

Murphy's Law, though frequently expressed humorously, is not entirely a joke. It is a recognition of the complexity of physical systems and a warning against over-simplifi-

cation and is comparable with Einstein's warning that "Everything should be made as simple as possible - but no simpler".¹

IMPORTANT COROLLARIES TO MURPHY'S LAW

- After it has worked successfully for two weeks it will fail during the first public demonstration.
- Equipment blows to protect fuses.
- Interchangeable parts aren't.
- Fail-safes don't.

Figure 11.2

This section of the seminar discusses the various physical effects which must be considered in the design of the hardware of mixed signal systems. Often such consideration will amount to a quick calculation to demonstrate that further consideration is not necessary, but sometimes extensive analysis,

or even actual experiments, will be necessary. However, the quick calculation must not be omitted, since problems are rarely obvious and often unexpected. The effects which must be considered will include many basic laws of physics.

BASIC LAWS INVOLVED IN THE DESIGN OF MIXED SIGNAL CIRCUITRY

- Ohm's Law
- Kirchoff's Law
- Faraday's Laws
- Lenz's Law

Figure 11.3

We therefore shall use as a section heading the major phenomenon considered in the section, but in the most general sense (for example, under "Resistance" we shall consider the non-ideal behavior of resistors, including noise, thermo- electric and inductive effects, which are not strictly issues of Ohm's Law).

When considering the effects of circuit conditions we are, of course, interested in

their effects on the performance of the system as a whole. Failure to allow for this is at the root of many of the problems which this section considers. For example, a 16-bit system divides its full-scale (FS) range into 2^{16} or 65536, which means that 1 LSB in a 10 V FS system is only 153 μ V. If we assume that we can tolerate errors of no more than 0.5 LSB, this calculation tells us that in a 16-bit system with 10 V FS we must keep the

total error to less than 76 μV , which is approximately equal to the thermoelectric voltage in a nichrome wirewound resistor with copper/nickel leads having about 2°C temperature difference between its ends.

Binary logic circuitry, on the other hand, has only two states, logic 0 and logic 1, and noise immunity of hundreds or thousands of millivolts. This is why circuit designers who have only worked with digital circuitry tend to overlook the sources of error which we are considering in this section of the seminar.

Figure 11.4 lists the sizes of 0.5 LSB at various resolutions (the values are given for

10 V fullscale since this is a classical converter range and where LSBs are given a mV value in this section of the seminar a 10 V FS is assumed unless explicitly stated otherwise - scaling to other values of FS is a trivial operation). Every analog designer should be familiar with this table, since not only does it allow the comparison of converters which are specified in different ways but it also indicates whether a design is reasonable or not - if noise or system errors amount to 1 mV there is little point in designing a system with more than 12-bits resolution.

BIT SIZES FOR 10 V FULLSCALE CONVERTERS

RESOLUTION	1 LSB	0.5 LSB	% FS	ppm FS	dB FS
4-bit	625mV	313mV	6.25	62500	-24
6-bit	156mV	78mV	1.56	15625	-36
8-bit	39mV	19.5mV	0.39	3906	-48
10-bit	9.76mV	4.88mV	0.098	977	-60
12-bit	2.44mV	1.22mV	0.024	244	-72
14-bit	610 μV	305 μV	0.0061	61	-84
16-bit	153 μV	76 μV	0.0015	15	-96
18-bit	38 μV	19 μV	0.0004	4	-108
20-bit	9.5 μV	4.8 μV	0.0001	1	-120
22-bit	2.4 μV	1.2 μV	0.000024	0.24	-132
24-bit	0.6 μV	0.3 μV	0.000006	0.06	-144

Figure 11.4.

RESISTANCE

RESISTANCE OF CONDUCTORS

Every engineer is familiar with resistors - little cylinders with wire ends - although perhaps fewer are aware of all their idiosyncrasies. Far too few engineers consider that all the wires and PC tracks with which their systems and circuits are assembled are also resistors.

At 25°C the resistivity of pure copper is 1.724E-6 ohm cm. The thickness of standard (1 ounce) PCB foil is 0.038 mm (0.0015"). The resistance of standard PCB copper is therefore 0.45 milliohms/square, which implies a resistance for the 0.25 mm track frequently used in computer designed digital circuitry of

18 milliohms/cm, which is quite large. Moreover the temperature coefficient of resistance for copper is about 0.4% /°C around room temperature, which can be a further inconvenience.

As an illustration of the effect of PCB track resistance consider a 16-bit ADC with a

5k ohm input resistance which has 5 cm of 0.25 mm PCB track between it and its signal source. This track has a resistance of approximately 0.09 ohms and introduces a gain error of 0.09 ohms / 5000 ohms (0.0018%) which is well over 1 LSB (0.0015% for 16 bits).

PRINTED CIRCUIT BOARD TRACK RESISTANCE

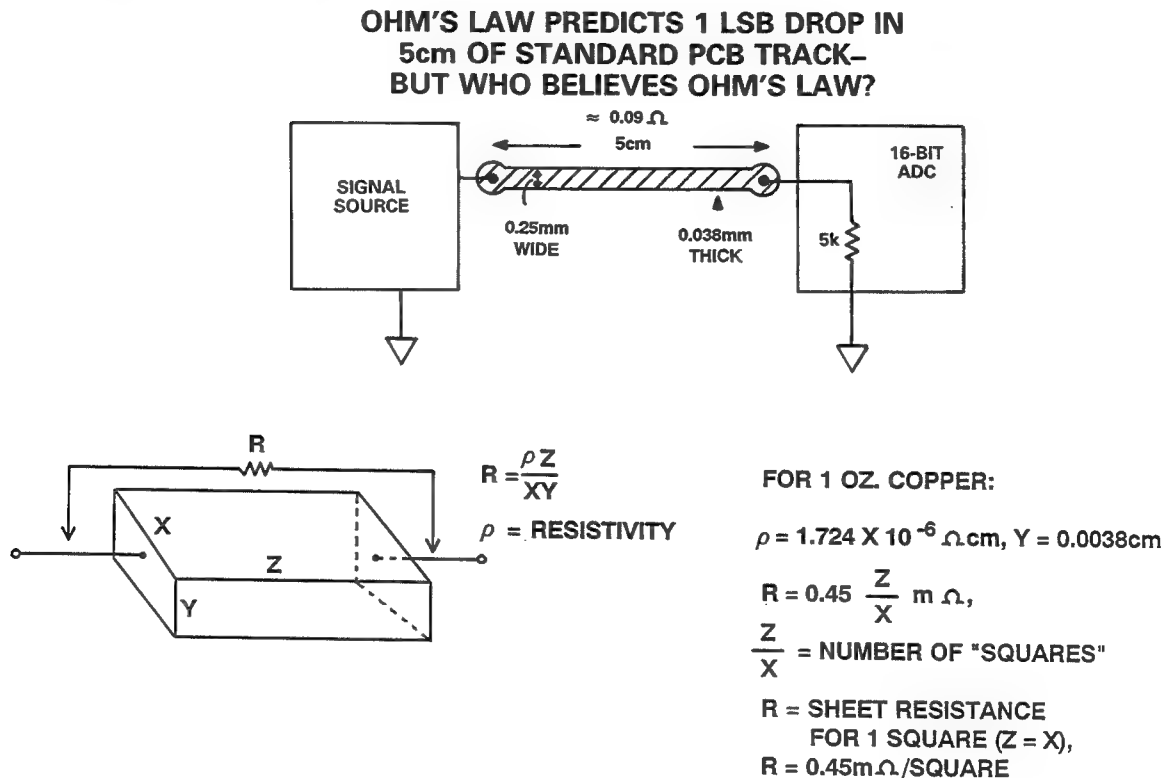


Figure 11.5

SKIN EFFECT

This, of course, is a DC effect. At high frequencies we must also consider the "skin effect" where inductive effects cause currents to flow only in the surface of conductors. This has the effect of increasing the resistance of a conductor at high frequencies (note that this effect is separate from the increase in impedance due to the effects of the self-inductance of conductors as frequency is increased - that

will be dealt with later). Skin effect is quite a complex phenomenon and detailed calculations are beyond the scope of this seminar. However a good approximation for copper is that the skin depth in centimeters is

$$\frac{6.6}{\sqrt{f}}, \text{ (f in Hz).}$$

Assuming that skin effects become important when the skin depth is less than 50% of the thickness of the PC foil this tells us that for normal 0.038 mm PC foil we must be concerned about skin effects at frequencies above approximately 12 MHz.


Where skin effect is important the resistance per square for copper is

$$2.6 \times 10^{-7} \sqrt{f} \text{ Ohms per square, (f in Hz)}$$

When calculating skin effects in PCBs it is important to remember that current flows in both sides of the PC foil (this is not necessarily the case in microstrip lines) so the resistance per square of PC foil is half the above value.

SKIN EFFECT

- HF Current flows only in thin surface layers



- Skin Depth: $6.61 / \sqrt{f}$ cm, f in Hz
- Skin Resistance: $2.6 \times 10^{-7} \sqrt{f}$ ohms per square, f in Hz
- Since skin currents flow in both sides of a PC track, the value of skin resistance in PCBs must take account of this

Figure 11.6

SKIN EFFECT

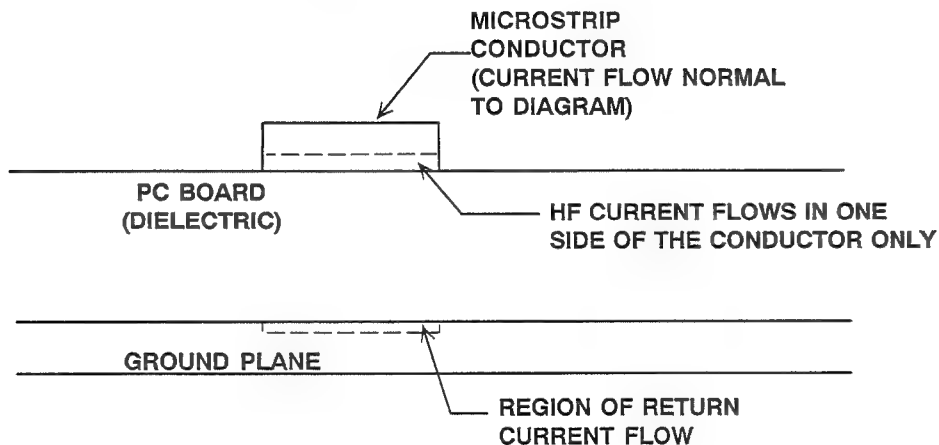


Figure 11.7

VOLTAGE DROP IN SIGNAL LEADS - "KELVIN" FEEDBACK

The gain error resulting from resistive voltage drop in signal leads is important only at high resolutions (as in the example) or where large signal currents flow. Where the load impedance is constant and resistive it

can be compensated by adjusting the overall system gain. In other circumstances it may often be removed by the use of "Kelvin" or "voltage sensing" feedback.


**USE OF A SENSE CONNECTION
MOVES ACCURACY TO THE LOAD**

Figure 11.8

Separate force and sense connections at a load remove any errors resulting from voltage drops in the force lead, but, of course, may only be used in systems where there is

negative feedback. It is also impossible to use such an arrangement to drive two or more loads with equal accuracy since feedback may only be taken from one point.

LEAKAGE IN INSULATORS

Just as conductors are improperly viewed as superconductors, so are insulators often mistakenly treated as perfect insulators, rather than very high resistances, which is the more accurate model.

Most printed circuit board materials are very good insulators, but they are not perfect, and inadequately cleaned PCB material may be quite a poor insulator. Furthermore, PCBs are anisotropic - even on a clean PCB different parts of the surface may have different resistivities, and the bulk resistance (between two plated through holes, for in-

stance) is generally lower than the surface resistance between two tracks.

Since the insulation resistance is so variable (and it will vary further with temperature and humidity) it is hard to predict in any particular circumstances but it is safe to assume that it is unlikely that the resistance between two conductors on a clean PCB will drop below 10^{10} - 10^{11} ohms, and with teflon PCB material (which is very expensive) will usually be over 10^{12} ohms.

GUARD RINGS

In applications where high impedances and very low currents are involved a guard ring may be used to minimize the effects of low insulation resistance. If critical high impedance nodes are surrounded by a ring of conductor which is at (or very close to) the potential of the node itself then the leakage current at the node will be minimized. If the node is at, or near to, ground then a grounded guard ring will be appropriate, if it is at some other potential it may be necessary to use a high input impedance buffer

amplifier, with its input connected to the node, to force the guard ring to the node potential. It is obvious that, in general, guard rings should be on both sides of the PCB with plated-through holes.

Nodes which are sufficiently sensitive to require guard rings should not contain plated through holes (unless the PCB is made of teflon) because, as mentioned above, the bulk resistivity of PCB material is less than the surface resistivity.

LEAKAGE RESISTANCE ON PCBs




Figure 11.9

An alternative to the use of a guard ring is to use teflon stand-off insulator(s) to support the high impedance point(s). If virgin teflon is used insulation resistance of around 10^{15} ohms is possible ("Virgin teflon" is a solid

piece of new teflon material which has been machined to shape and has not been welded together from powder or grains). The material of the rest of the circuit board need not have particularly high insulation resistance.

A VIRGIN TEFLON STANDOFF INSULATOR HAS MUCH LOWER LEAKAGE THAN A PCB TRACK




Figure 11.10

ELECTROSTATIC DAMAGE (ESD)

Where resistances are very high, especially in conditions of low humidity, there is always the possibility of electrostatic charge and electrostatic damage. A full discussion of electrostatic damage (ESD) and its prevention will be found in Analog Devices' Application Note on the subject, which is available free of charge from Analog Devices.²

This application note describes procedures to minimize the risk of electrostatic damage to sensitive devices. The basic principle of all ESD protection is to prevent a vulnerable item from being in the path of a discharge.

Many of the precautions used in factories are designed to minimize the possibility of any damaging discharge, even in the event of carelessness. When experienced engineers handle ICs they may dispense with most of the ESD protection apparatus and merely ensure that the IC is never in any potential discharge path: when taking a circuit from conductive foam, touch the foam to equalize charge before touching the circuit, similarly touch the foam with the hand before inserting the circuit in it, and hold your colleague's hand BEFORE passing the IC.

ELECTROSTATIC DISCHARGE (ESD)

ESD PREVENTION MANUAL



Figure 11.11

All integrated circuit structures are vulnerable to damage from the high voltages and high peak currents involved in even small electrostatic discharges but precision analog circuits suffer from a special disadvantage - the circuitry used to protect integrated circuit structures from ESD can often degrade the analog accuracy of the circuit where it is employed. Thus we have the choice between high performance and a high degree of protection. Which we choose will depend upon individual circumstances but it is essential to realize that the choice must be made - and if it is made in favour of accuracy then the circuit involved must not be exposed to electrostatic discharge.

A precision analog circuit exposed to ESD may not fail totally, but merely suffer degra-

dation of its analog performance, and possible reduction of life expectancy. When an IC is returned to Analog Devices for failure analysis of inadequate performance the first check that is made when the package is opened is a visual inspection for evidence of electrostatic damage - and this is found in a large percentage of cases.

An interesting example of an unobvious effect of ESD occurred in Finland, where very cold winters produce very low humidity and particularly severe electrostatic problems. A customer complained that the AD549 low bias current BIFET op-amp had poor long-term reliability and that its noise performance deteriorated over a few years of use.

ELECTROSTATIC DISCHARGE PROTECTION

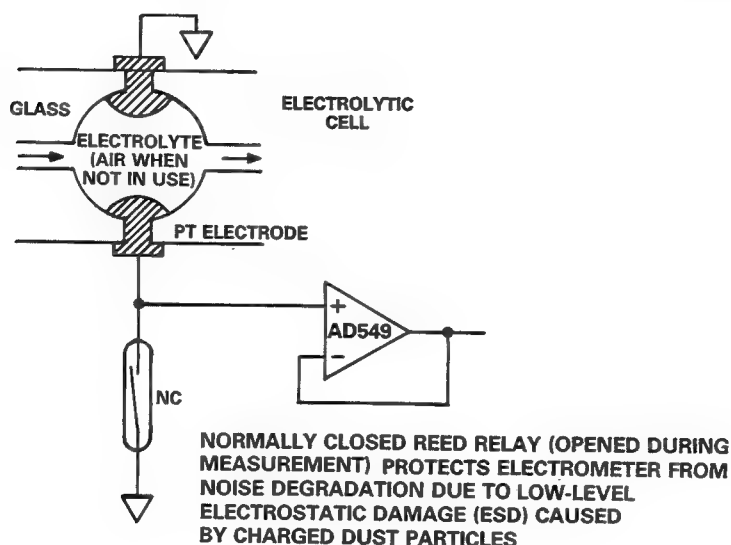


Figure 11.12

The amplifier was being used as a unity gain buffer with an electrochemical cell and the non-inverting input was connected to a platinum electrode and to nothing else. In use this electrode was immersed in electrolyte but after use it was washed (automatically) in deionized water and air dried. It was then left unconnected until the machine was next used.

Although there was no possibility of the electrode being touched at this time (it was

in the very center of the machine) it could encounter random particles of electrostatically charged dust - and the pulse currents as these dust particles discharged were sufficient to cause gradual deterioration of the noise figure. As soon as arrangements were made to ground the electrode when it was not in use (with an NC reed relay for minimum leakage) the problem disappeared.

PARASITIC EFFECTS IN RESISTORS

When we model a circuit, either informally or with a program such as SPICE, we generally assume that a resistor is a simple resistance. In fact any resistor is a much more

complex device containing, at the very least, an inductance, a noise source, a capacitor and two thermocouples.

THE EQUIVALENT CIRCUIT OF A RESISTOR IS NOT



BUT




Figure 11.13

Inductive resistors

All resistors have some inductance (as we shall see, a straight piece of wire has some inductance) but wirewound resistors actually consist of a coil of wire, which must inevitably be inductive. Even if the coil is “non-inductive” and consists of N clockwise turns and N anticlockwise turns there will still be some mismatch and residual inductance. Residual inductance values of up to $20\ \mu\text{H}$ can be expected in “non-inductive” wirewound resistors with values below $10\ \text{k}\ \Omega$, although above $10\ \text{k}\ \Omega$ the reactance of such a resistor is more likely to be a capacitance of around $5\ \text{pF}$.

Some film resistors are also inductive, consisting of a spiral of resistive material on a cylindrical ceramic body. Again values of a few μH are typical. High frequency circuits must not use inductive resistors since their impedance is not equal to their resistance and, indeed, varies with frequency. Even low frequency circuitry, where the inductance of the resistors would not seem to be a problem, may suffer from instability arising from unforeseen HF effects of resistor inductance (the transistors used in low frequency op-amps frequently have F_t of up to $1\ \text{GHz}$).

THERMO-ELECTRIC EFFECTS

Wirewound resistors have another problem. The junction of the resistance wire and the lead forms a thermocouple which has a thermoelectric EMF of $42 \mu\text{V}/^\circ\text{C}$ for the standard "Alloy 180"/Nichrome junction of an ordinary wirewound resistor. If a resistor is chosen with the [more expensive] copper/nichrome junction the value is $2.5 \mu\text{V}/^\circ\text{C}$. ("Alloy 180" is the standard component lead alloy of 77% copper and 23% nickel.)

Such thermocouple effects are unimportant at AC or where a resistor is at a uniform

temperature but if the dissipation in a resistor, or its location with respect to heat sources, can cause one of its ends to be warmer than the other then there will be a net thermoelectric EMF which will introduce a dc error into the circuit. With a normal wirewound resistor a temperature differential of only 4°C will introduce a dc error of $168 \mu\text{V}$ - which is greater than 1 LSB in a 10 V/16-bit system.

MINIMIZING THERMOCOUPLE EFFECTS IN WIREWOUND RESISTORS




Figure 11.14

The problem may be minimized by mounting wirewound resistors to ensure that temperature differentials are minimized. This may be done by ensuring that both leads are of equal length to equalize thermal conduction through them, by making any airflow (whether forced or natural convection) nor-


mal to the resistor body, and by taking care that both ends of the resistor are the same distance from any heat source on the PCB. Notwithstanding these precautions it is wiser to use resistors with copper, rather than "Alloy 180" leads, and to site them as far as possible from any heat source.

STABILITY & MATCHING

Thermal effects other than thermocouple effects will also affect the accuracy of circuits using resistors. Resistors are never completely stable with temperature and if either the temperature coefficients, or the actual temperatures of two resistors in a precision

circuit are mismatched then the performance of the circuit will suffer. Temperature mismatch of two identical resistors in similar environments may arise from differences in self-heating or other causes.³

GAIN OF 100 STAGE



- Resistor mismatch due to mismatch of temperature coefficients, mismatch of temperature (possibly due to self-heating), or both, can cause gain errors.
- Ideally, all resistors whose matching can affect accuracy should be fabricated on a single substrate.

Figure 11.15

Typical temperature coefficients of discrete resistors are apt to be around 100 ppm/°C or more. The best way to minimize the effects of resistor temperature coefficients and to eliminate the effects of different resistor temperatures is to ensure that all resistors whose resistor matching affects the accuracy of a system are built on a single substrate. This substrate may be the glass or ceramic substrate of a thin film resistor network.

A better alternative, when possible, is to use an integrated circuit having laser trimmed thin film resistors on the silicon substrate of the IC. The temperature coefficient of such resistors can be well below 20 ppm/°C, and the differential temperature coefficient between two resistors on the same substrate is of the order of 0.5 ppm/°C or less.

VOLTAGE VARIATION OF RESISTANCE

It is not possible to fabricate very high resistances on thin film or IC substrates, and high value discrete resistors are considerably less stable than lower value ones. It is inadvisable, therefore, to rely on the stability of high value resistors for the performance of a

system. Some types of high value resistor have another imperfection: they have a slightly non-linear voltage/current curve and do not obey Ohm's Law accurately.

HIGH VALUE RESISTORS

- Likely to be Less Stable

and

- Non-Linear With Voltage

Figure 11.16

JOHNSON NOISE

A final “imperfection” of resistors is an inconvenience but cannot properly be regarded as an imperfection as it is a fundamental

property of all resistors: thermal or Johnson noise.

RESISTOR JOHNSON NOISE

- All Resistors Have Noise: $V_n = \sqrt{4kTBR}$

T is Absolute Temperature

B is Bandwidth in Hertz

R is the Resistance in Ohms

k is Boltzmann’s Constant

(1.38E-23 J/K)

- It is possible to reduce the noise of a resistor by reducing T, B, or R but it is *NOT* possible to reduce k because Boltzmann is dead.

Figure 11.17

At any temperature above absolute zero all resistors have noise due to thermal motion of their structure. This noise, which is described by

$$V_n = \sqrt{4kTBR}$$

(Where k is Boltzmann’s constant: 1.38E-23 J/°K)

Johnson noise is present in ALL resistors and can only be reduced by reducing R, the resistance itself, B, the bandwidth of interest, or T, the temperature. Since the function involves a square root, the noise improvement for a drop in temperature from room temperature (298 °K) to liquid nitrogen (77 °K) is only of the order of 50%, so cooling a

resistor, unless liquid helium is involved, is unlikely to be very profitable.

Johnson noise is purely an effect of resistance. The Johnson noise of complex impedances consists only of the Johnson noise of

the resistive part of the impedance, so pure capacitance or inductance does not have Johnson noise, even though it has an impedance.

CAPACITANCE

STRAY CAPACITANCE

Where two conductors are not short-circuited together, or totally screened from each other by a conducting (Faraday) screen, there is a capacitance between them. There will therefore be a large number of capacitors associated with any circuit, which may or may not be considered in models of the circuit. Where high frequency performance matters (and even DC and VLF circuits may use devices with high Ft and therefore be vulnerable to HF instability) it is very important to consider the effects of stray capacity.

Any basic textbook will provide formulas for the capacitance of parallel wires, concentric spheres and cylinders, and many other configurations.⁴ The only example we need consider in this seminar is the parallel plate capacitor, which is formed by conductors on opposite sides of a PCB.

Neglecting edge effects, the capacitance of two parallel plates of area $A \text{ mm}^2$ and separation $d \text{ mm}$ in a medium of dielectric constant E_r relative to air is

$$0.00885 E_r A/d \text{ pF.}$$




Figure 11.18

From this formula we can calculate that for general purpose PCB material ($E_r = 4.7$, $d = 1.5$ mm) the capacitance between conductors on opposite sides of the board is just under $3\text{pF}/\text{cm}^2$. In general such capacitance will be parasitic, and circuits must be designed so that it does not affect their performance, but it is possible to use PCB ca-

pacitance in place of small discrete capacitors. However the dielectric properties of common PCB materials (teflon is an expensive exception) cause such capacitors to have a rather high temperature coefficient and to have poor Q at high frequencies, which makes them unsuitable for many applications.

CAPACITIVE NOISE & FARADAY SHIELDS

There is a capacitance between any two conductors separated by a dielectric (air or vacuum is a dielectric). If there is a change of

voltage on one there will be a movement of charge on the other. The basic model is shown in Figure 11.19.

CAPACITIVE COUPLING EQUIVALENT CIRCUIT





Figure 11.19

It is evident that the voltage coupled into Z_1 may be reduced by reducing the signal voltage, V_N , the frequency involved, the capacitance, or Z_1 , but frequently none of

these can be changed. The best solution is to insert a grounded conductor (known as a Faraday shield) between the noise source and the circuit which it affects.

CAPACITIVE SHIELDING

CAPACITIVE SHIELD INTERRUPTS THE COUPLING ELECTRIC FIELD



EQUIVALENT CIRCUIT ILLUSTRATES HOW A CAPACITIVE SHIELD CAUSES THE NOISE CURRENTS TO RETURN TO THEIR SOURCE WITHOUT FLOWING THROUGH Z_1




Figure 11.20

The Faraday shield is easily implemented and almost invariably successful. For this reason capacitively coupled noise is rarely an intractable problem. However, to be effective the shield must completely block the electric field between the noise source and the shielded circuit and must be connected so that the noise current returns to its source without flowing in any part of the circuit where it might introduce conducted noise. A conductor intended as a Faraday shield must never be left unconnected as this almost


always increases capacity and exacerbates the problem.

An example of this problem is seen in sidebraced ceramic IC packages. These DIP packages have a small square conducting kovar lid soldered onto a metallized rim on the ceramic package top. Package manufacturers offer only two options: the metallized rim may be connected to one of the corner pins of the package or it may be left unconnected. Most logic circuits have a ground pin at one of the package corners and therefore

the lid is grounded. Many analog circuits do not have a ground pin at a package corner and the lid is left floating - such circuits turn out to be far more vulnerable to electric field

noise than the same chip in a plastic DIP package where the chip is completely unshielded.

CAPACITIVE EFFECTS DUE TO METAL LIDS



- **SIDEBRAZE CERAMIC D.I.L. PACKAGES SOMETIMES HAVE ISOLATED METAL LIDS**
- **THESE ARE VULNERABLE TO CAPACITIVE INTERFERENCE AND SHOULD BE GROUNDED (IF POSSIBLE)**

Figure 11.21

Whatever the environmental noise level, it is good practice for the user to ground the lid of any sidebrazed ceramic IC where the lid is not grounded by the manufacturer - this can be done with a wire soldered to the lid (this will not damage the device as the chip is thermally and electrically isolated from the lid). If soldering to the lid is unacceptable a grounded phosphor-bronze clip may be used to make the ground connection, or conductive

paint from the lid to the ground pin. Never attempt to ground such a lid without verifying that it is, in fact, unconnected, as occasionally device types will be found with the lid connected to a power supply rather than to ground!

One case where a Faraday shield is impracticable is between the bondwires of an integrated circuit chip. This has important consequences.

STRAY CAPACITY BETWEEN CHIP BONDWIRES




Figure 11.22

The stray capacitance between two chip bondwires and their associated leadframes is of the order of 0.2 pF. (Note this is “of the order” NOT “of the close order” - observed values generally lie between 0.05 and 0.6 pF.) If we have a high resolution converter (ADC or DAC) which is connected to a high speed data bus then each line of the data

bus, which will be carrying noise with 2-5 V/ns dV/dT , is connected to the converter analog port via this stray capacitance. Whenever the bus is active intolerable amounts of noise will be capacitively coupled to the analog port and will seriously degrade the performance of which the converter is capable.

WITH A HIGH PERFORMANCE CONVERTER
ON A HIGH SPEED DATA BUS, IT IS
NOT POSSIBLE
TO SHIELD THE ANALOG PORT
FROM THE DIGITAL NOISE





Figure 11.23

Present technology offers no cure for this problem, which also limits the performance possible from broadband monolithic mixed signal ICs having analog and digital circuitry on a single chip. However, it may be avoided quite simply by not connecting the databus directly to the converter but by using a

latched buffer as an interface. This solution costs money, occupies board area, reduces reliability (very slightly), consumes power and complicates design - but it does improve the signal-to-noise ratio of the converter. The designer must decide whether it is worthwhile in individual cases.

BUFFER LATCH USED AS FARADAY SHIELD



- A BUFFER/LATCH CAN ACT AS A FARADAY SHIELD BETWEEN A FAST DATA BUS AND A HIGH PERFORMANCE CONVERTER
- IT ADDS COST, BOARD AREA, POWER CONSUMPTION, RELIABILITY REDUCTION, DESIGN COMPLEXITY AND IMPROVED PERFORMANCE

Figure 11.24

PARASITIC EFFECTS IN CAPACITORS

Just as we are too willing to assume that a resistor is a perfect resistor, so do we underestimate the parasitic components associated with a capacitor. Fig 25 shows the

ideal, the general model of a real capacitor, and the simplified models which are adequate for the analysis of non-ideal behavior in most applications.

EQUIVALENT CIRCUITS OF A REAL CAPACITOR




Figure 11.25

Capacitors are used for coupling (passing AC signals while blocking DC), for decoupling (removing AC superimposed on DC in both power and signal circuitry), for building

filters or frequency-selective networks, and for storing charge in "sample and hold" circuits (also known as "track and hold" circuits or SHAs, SAHs or THAs).

CAPACITOR LEAKAGE

In coupling and SHA applications the leakage of the capacitor can be important. Electrolytic capacitors, where the dielectric is formed by an electrochemical reaction, have relatively high leakage currents of microamperes or even more and so are not used in applications where leakage matters. The leakage of electrolytic capacitors is greater during the first few minutes of operation after a period of storage (the leakage current while the capacitor is in use keeps the dielectric in good condition and it may deteriorate slightly in storage) - this feature can be important in equipment which must perform correctly after a long quiescent period.

The leakage of tantalum electrolytic capacitors is lower than that of aluminium ones and so in applications where capacitances of tens of microfarads or more (which can be easily achieved only with electrolytic

capacitors) are required tantalum ones are used, despite their extra cost, if particularly low values of leakage current are necessary. At room temperature the leakage of aluminium electrolytic capacitors is of the order of 20 nA/ μ F and that of tantalum ones is 5 nA/ μ F.

Another feature of electrolytic capacitors, both aluminium and tantalum, is that most of them are polarized and require a DC bias for correct operation - a reverse bias may do damage and will certainly increase leakage (unpolarized electrolytic capacitors, which may be biased in either direction, do exist but they are uncommon, and considerably larger than the polarized variety).

Most other types of capacitor have leakage resistances in excess of hundreds of gigohms so that for most applications their leakage currents can be disregarded.

SERIES/LOSS RESISTANCE

The series resistance of capacitors causes them to dissipate power when high AC currents are flowing in them. This can have serious consequences at RF and in supply decoupling capacitors carrying high ripple

currents but is unlikely to have much effect in precision analog circuitry. The series inductance, however, can have very inconvenient consequences.

INDUCTANCE OF CAPACITORS

The transistors used in precision analog circuits have transition frequencies (F_t) of hundreds of MHz or even several GHz, even though the precision circuitry itself may be operating at DC or low frequencies. This makes it essential that the power supply terminals of such circuits should be decoupled properly at high frequency.

A common structure for capacitors is two sheets of metal foil separated by sheets of plastic or paper dielectric and formed into a roll. Such a structure has considerable self inductance and behaves as an inductance rather than a capacitor at frequencies of more than a few MHz. It is therefore inadvisable to use electrolytic, paper or plastic film

capacitors for decoupling at high frequencies.

Monolithic ceramic capacitors have very low series inductance (they are formed of a multilayer sandwich of metal films and ceramic dielectric and all the films are joined to a bus-bar rather than being connected in series). They are therefore ideal for high frequency decoupling. However, monolithic ceramic capacitors can be microphonic, and some types may be self-resonant with comparatively high Q . Disc ceramic capacitors, on the other hand, are sometime quite inductive, although less expensive.

The best way of ensuring that an analog circuit is adequately decoupled at both high and low frequencies is to use a tantalum bead capacitor in parallel with a monolithic ceramic one. The combination will have high capacitance but will remain capacitive at VHF frequencies. It is generally unnecessary to have a tantalum capacitor on each individ-

ual IC, if there is less than 10 cm of reasonably wide PC track between each IC and the tantalum capacitor it is possible to share one tantalum capacitor among several ICs.

There is little point in taking great care in the choice of a non-inductive capacitor if it is then unsuitably mounted. Short lengths of wire have appreciable inductance so HF decoupling capacitors must be mounted as close as possible to the points that they are decoupling with short, wide PC tracks. Ideally HF decoupling capacitors should be surface-mount parts to eliminate lead inductance, but wire-ended capacitors are permissible provided the device leads are no longer than 1.5 mm. It is also important to understand where HF decoupling currents should flow and why HF decoupling is more important at some points than at others - the subject is covered at some length in an Analog Devices Application Note.⁵

HIGH FREQUENCY DECOUPLING (REQUIRED EVEN BY LF ANALOG CIRCUITS)

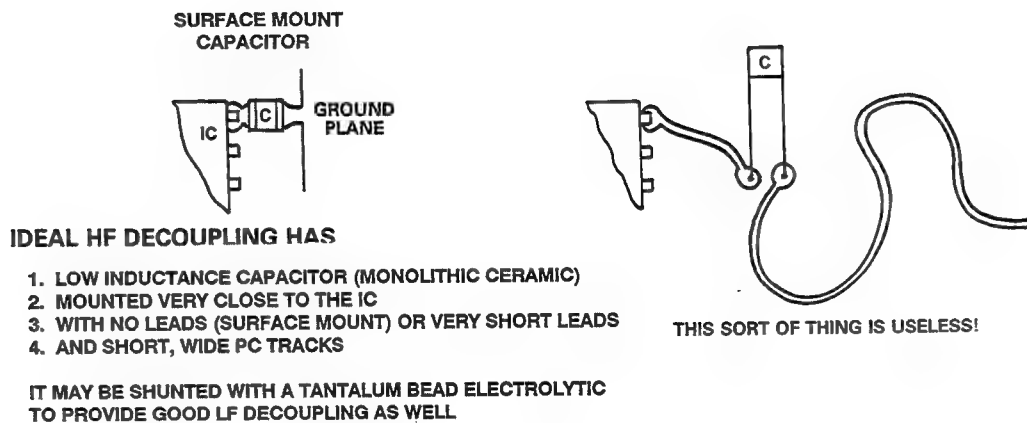


Figure 11.26

HF instability in analog circuits is more common than is realized. Oscillation at hundreds of MHz will cause serious malfunction of precision circuitry but may not affect an oscilloscope (indeed the presence of an oscilloscope probe may damp the oscillation, so that the circuit works only when an oscilloscope is attached to it - this is an impor-


tant diagnostic clue). It is quite good practice to use a broadband spectrum analyzer (say 1-1500 MHz) and a low capacity FET probe to check for parasitic oscillation any analog circuit which is malfunctioning for no obvious reason. This test will also show if the malfunction is due to the presence of a strong RF field from an external source.

DIELECTRIC ABSORPTION

Monolithic ceramic capacitors are excellent for HF decoupling but they have considerable dielectric absorption, which makes them unsuitable for use as the hold capacitor of an SHA. Dielectric absorption causes a capacitor which is quickly discharged and

then open-circuited to recover some of its charge. Since the amount of charge recovered is a function of its previous charge this is, in effect, a charge memory and will cause errors in any SHA where dielectric absorption is present in the hold capacitor.

CAPACITORS HAVING SIGNIFICANT D.A. ARE USELESS FOR SAMPLE AND HOLD APPLICATIONS



DIELECTRIC ABSORPTION CAUSES A BRIEFLY DISCHARGED CAPACITOR TO RECOVER A PERCENTAGE OF ITS PREVIOUS CHARGE ON BEING OPEN CIRCUITED.

Figure 11.27

Capacitors for this application should therefore be selected to have minimal dielectric absorption. The best strategy is to use a SHA which is supplied with an internal capacitor or where the SHA manufacturer supplies the capacitor with the SHA. If this is not possible (sometimes one may require a longer hold time - and hence extra capacity) a capacitor should be chosen which has its low dielectric absorption (DA) specified on its

data sheet.

Such capacitors are normally plastic dielectric types (polystyrene, polypropylene or teflon) but it is not safe to use just any plastic dielectric capacitor with a SHA as special processing and testing is necessary to ensure that it has low DA. For use with a SHA a capacitor should be chosen which is specified for low DA applications.

INDUCTANCE

STRAY INDUCTANCE

All conductors are inductive and at high frequencies the inductance of even quite short pieces of wire may be important. The inductance of a straight wire of length L mm and circular cross-section with radius R mm in free space is

$$0.0002L \left[\ln \left(\frac{2L}{R} \right) - 0.75 \right] \mu\text{H}.$$

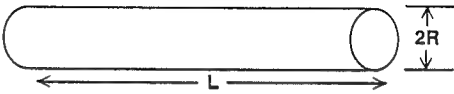
The inductance of a strip conductor (an approximation to a PC track) of width W mm and thickness H mm in free space is

$$0.0002L \left[\ln \left(\frac{2L}{W+H} \right) + 0.2235 \left(\frac{W+H}{L} \right) + 0.5 \right] \mu\text{H}.$$

In real systems these formulas both turn out to be approximate but they do give some idea of the order of magnitude of inductance involved. They tell us that 1 cm of 0.5 mm o.d. wire has an inductance of 7.26 nH and 1 cm of 0.25 mm PC track has an inductance of 9.59 nH - these figures are reasonably close to measured results.

At 10 MHz an inductance of 7.26 nH has an impedance of 0.46 ohm and so can give rise to 1% error in a 50 ohm system.


INDUCTANCE



L, R in mm

$$\text{WIRE INDUCTANCE} = 0.0002L \left[\ln \left(\frac{2L}{R} \right) - 0.75 \right] \mu\text{H}$$

EXAMPLE: 1cm of 0.5mm o.d. wire has an Inductance of 7.26nH
($2R = 0.5\text{mm}$, $L = 1\text{cm}$)



$$\text{STRIP INDUCTANCE} = 0.0002L \left[\ln \left(\frac{2L}{W+H} \right) + 0.2235 \left(\frac{W+H}{L} \right) + 0.5 \right] \mu\text{H}$$

EXAMPLE: 1cm of 0.25 mm PC track has an Inductance of 9.59 nH
($H = 0.038\text{mm}$, $W = 0.25\text{mm}$, $L = 1\text{cm}$)

Figure 11.28

MUTUAL INDUCTANCE

Another consideration regarding inductance is the separation of outward and return currents. As we shall discuss in more detail later, Kirchoff's Law tells us that current flows in closed paths - there is always an outward and return path. The whole path forms a single-turn inductor. If the area

enclosed by the turn is large the inductance, and hence the AC impedance, will also be large, whereas if the outward and return paths are close together the inductance will be much smaller. The principle is illustrated in Fig 11.29.

NONIDEAL AND IMPROVED SIGNAL ROUTING




Figure 11.29

The nonideal routing in Figure 11.29 has another drawback - the large area enclosed by the conductor produces extensive external magnetic fields, which may interact with other circuits and cause unwanted coupling. Similarly the large area is more vulnerable

to interaction with external magnetic fields, which can induce unwanted signals in the loop. The basic principle is illustrated in Figure 11.30 and is a common mechanism for the transfer of unwanted signals (noise) between circuits.

BASIC PRINCIPLES OF INDUCTIVE COUPLING

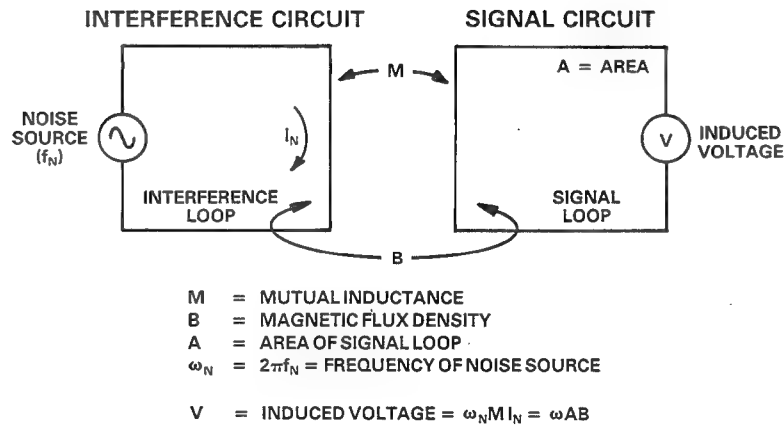


Figure 11.30

As with most other sources of noise, as soon as we define the principle at work we can see ways of reducing the effect. In this case reducing any or all of the terms in the equations in Figure 11.30 will reduce the coupling. Reducing the frequency or amplitude of the current causing the interference

may be impracticable but it is frequently possible to reduce the mutual inductance between the interfering and interfered with circuits by reducing loop areas on both sides and, possibly, increasing the distance between them.

PROPER SIGNAL ROUTING REDUCES MUTUAL INDUCTANCE





Figure 11.31


Mutual inductance is a common problem in ribbon cables, especially when a single return is common to several signal circuits. Separate signal and return lines for each

signal circuit reduces the problem, and using a cable with twisted pairs for each signal circuit is even better (but more expensive and often unnecessary).


MUTUAL INDUCTANCE AND SIGNAL COUPLING IN RIBBON CABLE



**FLAT RIBBON CABLE WITH SINGLE
RETURN HAS LARGE MUTUAL
INDUCTANCE BETWEEN CIRCUITS**



**SEPARATE AND ALTERNATE SIGNAL
AND RETURN LINES FOR EACH CIRCUIT
REDUCE MUTUAL INDUCTANCE**




**TWISTED PAIRS REDUCE MUTUAL
INDUCTANCE STILL FURTHER**

Figure 11.32

Shielding magnetic fields to reduce mutual inductance is sometimes possible but is by no means as easy as shielding electric fields with a Faraday shield. HF magnetic fields are blocked by conductive material, while LF and DC fields may be screened by a

shield made of mu-metal sheet. Mu-metal is an alloy having very high permeability, but it is expensive, its magnetic properties are damaged by mechanical stress, and it will saturate if exposed to too high fields. Its use, therefore, should be avoided where possible.

MAGNETIC SHIELDING



- Magnetic shielding is not as easily accomplished as electrostatic shielding, but may be done at HF with a simple conducting screen, and at LF and DC with a screen of high permeability material such as Mu-metal.

Figure 11.33


RINGING

An inductor in series or parallel with a capacitor forms a resonant, or “tuned”, circuit, whose key feature is that it shows marked change in impedance over a small range of frequency (how sharp the effect is depends on the Q of the tuned circuit). The effect is widely used to define the frequency response of narrow-band circuitry but can also be a source of problems.

If stray inductance and capacitance (which may or may not be stray) in a circuit


should form a tuned circuit then that tuned circuit may be excited by signals in the circuit and ring at its resonant frequency. A common example is shown in Figure 11.34 where the resonant circuit formed by an inductive power line and its decoupling capacitor may be excited by pulse currents drawn by the IC.

RESONANT CIRCUITS FORMED BY DECOUPLED POWER LINES



EQUIVALENT CIRCUIT
OF DECOUPLED POWER
LINE - RESONANT AT

$$f = \frac{1}{2\pi\sqrt{LC}}$$



SMALL SERIES RESISTANCE
CLOSE TO THE IC REDUCES THE Q

Figure 11.34

The effect may be minimized by lowering the Q of the inductance, which is most easily

done by inserting a small resistance in the power line, close to the IC.

PARASITIC EFFECTS IN INDUCTORS

Although inductance is one of the fundamental properties of an electronic circuit, inductors are less common as precision components than resistors and capacitors. This is because they are harder to manufacture, less stable, and less physically robust than resistors and capacitors. It is relatively easy to manufacture stable precision inductors with inductances from nH to tens or hundreds of μH , but larger valued devices tend to be less stable, and large.

As we might expect in these circumstances, circuits are designed, where possible, to avoid the use of precision inductors. We find that stable precision inductors are relatively rarely used in precision analog

circuitry, except in tuned circuits for high frequency narrow band applications.

Of course they are widely used in power filters, switching power supplies and other applications where lack of precision is unimportant. The important features of inductors used in such applications are their current carrying and saturation characteristics, and their Q. If an inductor consists of a coil of wire with an air core its inductance will be essentially unaffected by the current it is carrying, but if it is wound on a core of a magnetic material (magnetic alloy or ferrite) its inductance will be non-linear since at high currents the core will start to saturate.

SATURATION

- Inductors with solid cores (magnetic alloy or ferrite)
 - will behave non-linearly
 - if required to carry too much current.
- This is unlikely to be a direct problem in precision circuitry but may affect power supply noise performance and thus affect precision circuitry indirectly.

Figure 11.35

Such saturation will reduce the efficiency of the circuitry employing the inductor and is liable to increase noise and harmonic generation.

As mentioned above, inductors and capacitors together form tuned circuits. Since all in-

ductors will have some stray capacity, all inductors will have a resonant frequency (which will normally be published on their data sheet) and should only be used as precision inductors at frequencies well below this.

STRAY CAPACITANCE MAKES ALL INDUCTORS INTO TUNED CIRCUITS




Figure 11.36

Q OR “QUALITY FACTOR”

The other parasitic characteristic of inductors is their Q (or “Quality Factor”), which is the ratio of their reactive impedance to their resistance.

$$Q = 2 \pi f L / R$$

It is rarely possible to calculate the Q of an inductor from its DC resistance since skin effect (and core losses if the inductor has a magnetic core) ensure that the Q of an induc-

tor at high frequencies is always lower than that predicted from DC values.

Q is also a characteristic of tuned circuits (and of capacitors - but capacitors generally have sufficiently high values of Q that it may be disregarded for most practical purposes). The Q of a tuned circuit, which is generally very similar to the Q of its inductor (unless it is deliberately lowered by the use of an additional resistor), is a measure of its bandwidth around resonance.

Q or “QUALITY FACTOR”

- The Q of an inductor or resonant circuit is a measure of the ratio of its reactance to its resistance.

$$Q = 2 \pi f L / R$$

- The resistance is the HF and *NOT* the DC Value.
- The 3dB bandwidth of a single tuned circuit is F_c/Q where F_c is the center frequency.

Figure 11.37

LC tuned circuits rarely have Q of much more than 100 (3 dB bandwidth of 1%) but ceramic resonators may have Q of thousands

and quartz crystals have Q of tens of thousands.


GROUNDING & SIGNAL ROUTING

SIGNAL RETURN CURRENTS

Kirchoff's Law tells us that at any point in a circuit the algebraic sum of the currents is zero. This tells us that all currents flow in

circles and, particularly, that the return current must always be considered when analyzing a circuit.⁶

KIRCHOFF'S LAW



AT ANY POINT IN A CIRCUIT
THE ALGEBRAIC SUM OF THE CURRENTS IS ZERO
OR
WHAT GOES OUT MUST COME BACK
WHICH LEADS TO THE CONCLUSION THAT
ALL VOLTAGES ARE DIFFERENTIAL
(EVEN IF THEY'RE GROUNDING)

Figure 11.38

Most people consider the return current when considering a fully differential circuit, but when considering the more usual circuit where a signal is referred to "ground" it is

common to assume that all the points on the circuit diagram where the ground symbol is to be found are at the same potential. This is unwise.

THE IDEAL GROUND




Figure 11.39

GROUND NOISE & GROUND LOOPS

A more realistic model of ground is shown in Figure 11.40. Not only does the return current flow in the complex impedance which exists between the two "ground" points shown in Figure 11.39, giving rise to a volt-

age drop in the total signal path, but external currents may also flow in the same path, generating uncorrelated noise voltages which are seen by the ADC.

A MORE REALISTIC GROUND

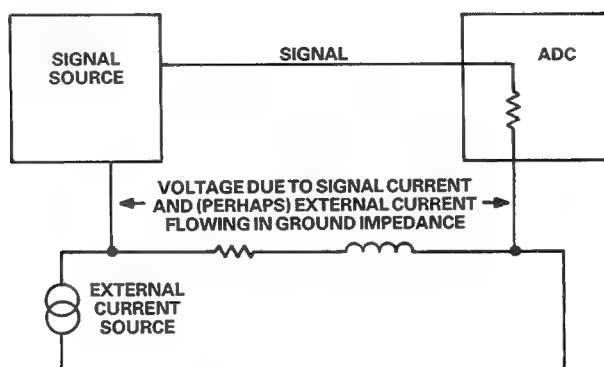


Figure 11.40

It is evident, of course, that other currents can only flow in the ground impedance if there is a current path for them. Figure 11.40 shows such a path at "ground" potential, which is the notorious "Ground Loop",

but equally severe problems could be caused by a circuit sharing an unlooped ground return with the signal source but drawing a large and varying current from its supply and ground return.

ANY CURRENT FLOWING IN A COMMON GROUND MAKES NOISE; A GROUND LOOP IS NOT NECESSARY

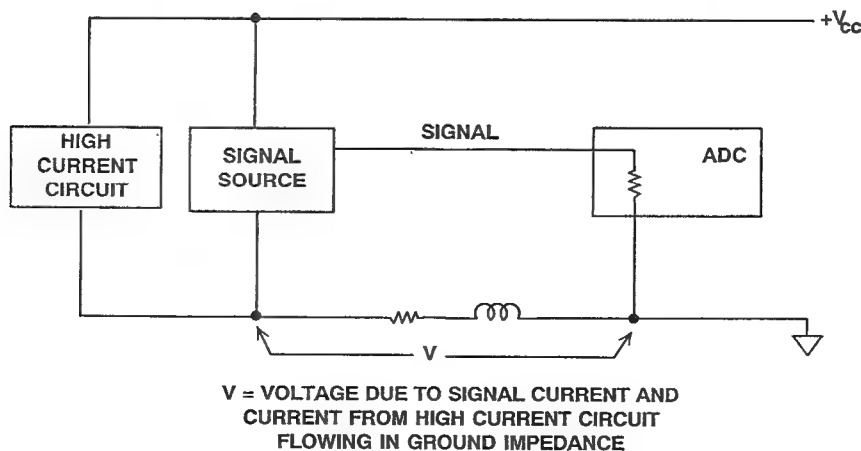


Figure 11.41

It is evident from Figure 42 that if a ground network contains loops there is a greater danger of it being vulnerable to EMF's induced by external magnetic fields,

and of ground current "escaping" from high current areas to cause noise in sensitive regions. For these reasons ground loops are best avoided.

GROUND LOOP

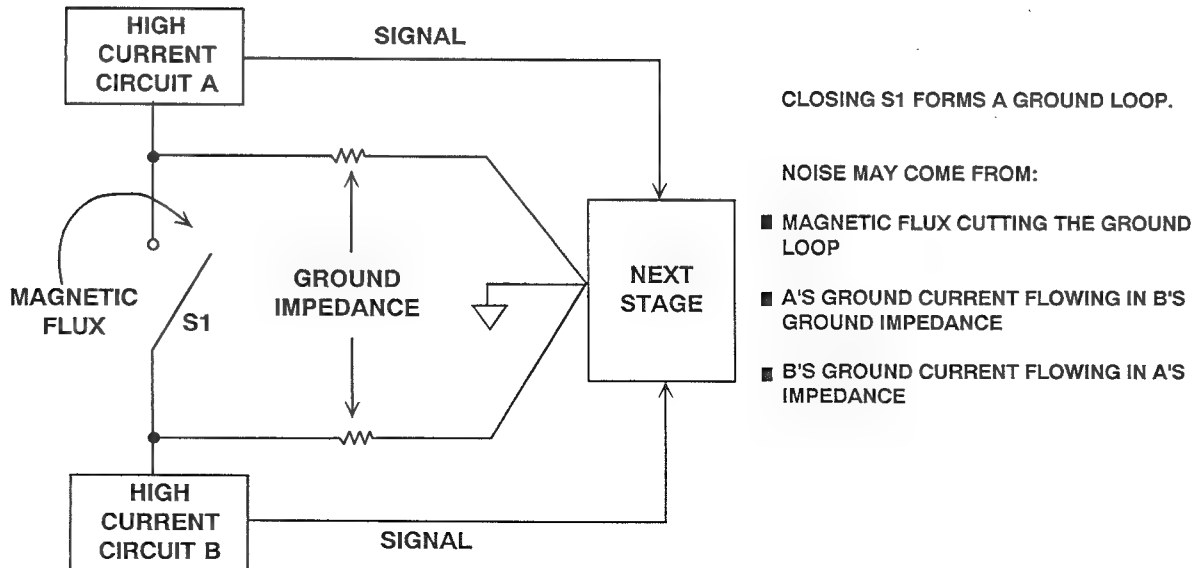


Figure 11.42

However, there are situations where looped grounds are unlikely to cause unacceptable noise and the configuration may actually offer benefits in the form of safety or reduced impedance. In such circumstances the optimum ground arrangement may contain loops. Sensible engineers should not allow the almost superstitious dread inspired by the term "ground loop" to prevent the adoption of such designs, if careful analysis and experiment has shown that they actually are optimum.

There are a number of possible ways of attacking the problem of ground noise, apart from the (presently) impracticable one of using superconducting grounds. It is rare for a single method to be used to the exclusion of all others, and systems generally contain a mixture of approaches. For the purposes of description, however, it is better to describe each approach separately.

STAR (MECCA) GROUNDS

The “star” or “Mecca” ground philosophy builds on the theory that there is a single point in a circuit to which all voltages are

referred. This is known as the “star” or “Mecca” point.

STAR (MECCA) GROUNDS

If all signal voltages in a system are measured with respect to a single point, that point is said to be the *star* ground of the system.

Figure 11.43

This philosophy is reasonable but frequently encounters practical difficulties. For example if we design a system with a star ground, drawing all the signal paths to minimize signal interaction and the effects of high impedance signal or ground paths, we frequently find, when the power supplies are added to the circuit diagram, that the power supplies either add unwanted ground paths

or that supply currents, flowing in existing ground paths, are sufficiently large, or noisy, or both, as to corrupt the signal transmission. This problem may often be avoided by having separate power supplies for different parts of the circuit - separate analog and digital supplies, and separate analog and digital grounds joined at the star point, are common in mixed signal applications.

SEPARATE ANALOG AND DIGITAL GROUNDS

Digital circuitry is noisy. Saturating logic draws large fast current spikes from its supply during switching and, having noise immunity of hundreds of millivolts or more, has little need of high levels of supply decoupling.

Analog circuitry, on the other hand, is very vulnerable to noise in supplies or grounds. It is therefore sensible to separate analog and digital circuitry to prevent digital noise from corrupting analog performance. Such separation will involve separation of

both grounds and power supplies, which may be inconvenient in a mixed signal system. Nevertheless, if a system is to give the full performance of which it is capable it is often essential to have separate analog and digital grounds and power supplies. The fact that some analog circuitry will operate from a single +5 V supply does NOT mean that it may safely be operated from the same noisy +5 V supply as the microprocessor and dynamic RAM, the electric fan, and the solenoid jackhammer!

SUPPLY & GROUND NOISE

- Digital circuitry is noisy
- Analog circuitry is quiet
- Circuit noise from digital circuitry carried by power and ground leads can corrupt precision analog circuitry
- It is advisable to separate the power and ground of the digital and analog parts of a system
- Analog and digital grounds must be joined at ONE point

Figure 11.44

However, analog and digital ground in a system must be joined at some point to allow signals to be referred to a common potential. This star point, or analog/digital common point, is chosen so that it does not introduce digital currents into the ground of the analog part of the system - it is often convenient to make the connection at the power supplies.

Many ADCs and DACs have separate “analog ground” and “digital ground” pins,

and users are advised, on the data sheets, to connect these pins together at the device package. This seems to conflict with the advice to connect analog and digital ground at the power supplies, and, in systems with more than one converter, with the advice to join the analog and digital ground at a single point.

ANALOG GROUND & DIGITAL GROUND

- Monolithic & hybrid ADCs frequently have separate AGnd & DGnd pins which must be joined together at the device.
- This is not done from a desire to be difficult, but because the voltage drop in the bondwires is too large to allow the connection to be made internally.
- The best solution to the grounding problem arising from this requirement is to connect both pins to system “analog ground”.
- It is likely that neither the digital noise so introduced in the system AGnd, nor the loss of digital noise immunity, will seriously affect the system performance.

Figure 11.45

There is, in fact, no conflict. The labels "analog ground" and "digital ground" on these pins refer to the parts of the converter to which the pins are connected, and not to the system grounds to which they must go. In general these two pins should be joined together and to the analog ground of the system. It is not possible to join the two pins within the IC package because the analog part of the converter cannot tolerate the

voltage resulting from the digital current flowing in the bond wire to the chip.

If these pins are connected in this way the digital noise immunity of the converter is diminished by the amount of common-mode noise between the digital and analog system grounds. Since digital noise immunity is of the order of hundreds or thousands of millivolts this is unlikely to be important.

ANALOG GROUND (AGND) AND DIGITAL GROUND (DGND) OF ADCs/DACs SHOULD BE RETURNED TO SYSTEM ANALOG GROUND

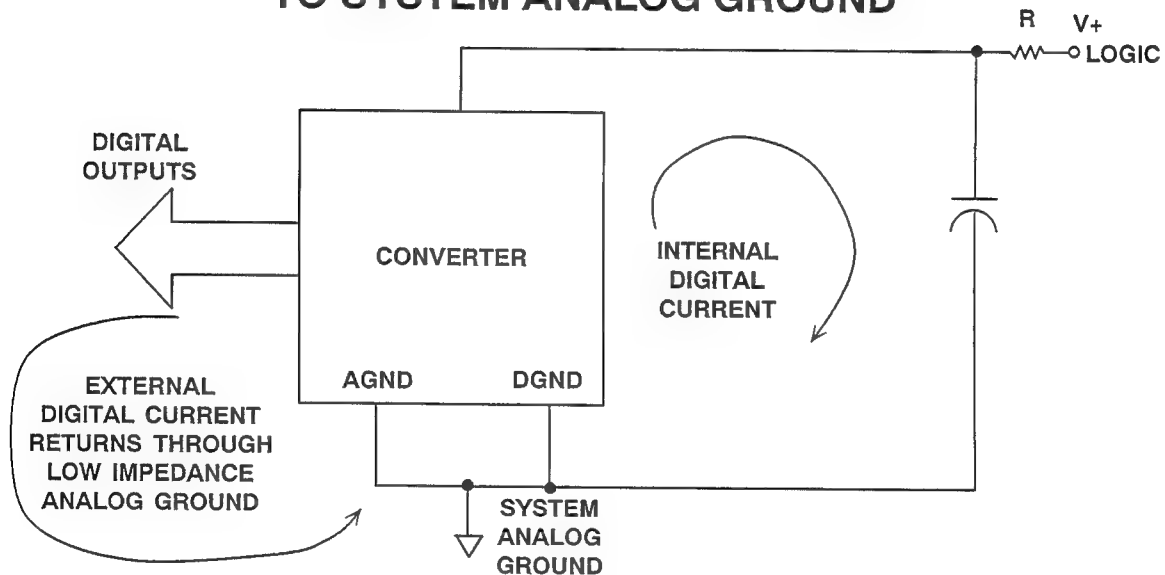


Figure 11.46

The analog noise immunity is diminished only by the external digital currents of the converter itself flowing in the analog ground. These currents should be quite small, and can be minimized by ensuring that the converter outputs do not drive large fanouts. If the logic supply to the converter is isolated with a small resistance and decoupled to analog ground with a $0.1\mu\text{F}$ capacitor sited

as close to the converter as possible all the internal digital currents of the converter will return to ground through the capacitor and will not appear in the external ground circuit. If the analog ground impedance is as low as it should be for adequate analog performance the additional noise due to the external digital ground current should rarely present a problem.

GROUND PLANES

Related to the star ground system is the use of a ground plane. One side of a double-sided PCB, or one layer of a multi-layer one, is made of continuous metal, which is used as

ground. The theory behind this is that the large amount of metal will have low resistance and as low inductance as is possible.

GROUND PLANES

- One entire side or layer of a PCB is continuous grounded conductor.
- This gives minimum ground resistance and inductance but is not always sufficient to solve all grounding problems.
- Breaks in ground planes can improve or degrade circuit performance - there is no general rule.
- Twenty years ago ground planes were difficult to fabricate. Today they are not.
- If your PCB facility objects to fabricating ground planes - GET A NEW PCB FACILITY!

Figure 11.47

It is sometimes argued that ground planes should not be used because they are liable to introduce problems in manufacture and assembly. Such an argument may have had a limited validity twenty years ago when PCB adhesives were less well developed, wave-soldering less reliable, and solder resist techniques less well understood, but today it should not be tolerated.

Ground planes solve many ground impedance problems, but not all. Even a continuous sheet of copper foil has residual resistance and inductance and in some circumstances they can be enough to prevent proper circuit function. Figure 11.48 shows such a problem - and a possible solution.

A SLIT IN A GROUND PLANE CAN RECONFIGURE CURRENT FLOW FOR BETTER ACCURACY




Figure 11.48

Consider a ground-plane PCB 100 mm wide with a ground connection at one end and a power amplifier at the other drawing 15A. If the ground plane is 0.038 mm thick and 15 A flows in it there will be a voltage drop of $68 \mu\text{V}/\text{mm}$. This voltage drop would cause quite serious problems to any ground-

referenced precision circuitry sharing the PCB. However, if we slit the ground plane so that high current does not flow in the region of the precision circuitry we can possibly solve the problem - even though the voltage gradient will increase in those parts of the ground plane where the current does flow.

TRANSMISSION LINES

A break in a ground plane is not always a good thing. We earlier considered the benefits of outward and return signal paths being close together so that inductance is minimized. As we saw in Figure 11.7, when an HF signal flows in a PC track running over a ground plane the arrangement functions as a microstrip transmission line and the majority of the return current flows in the ground plane underneath the line.

The characteristic impedance of the line will depend upon the width of the track and the thickness and dielectric constant of the

PCB material. For most lower frequency applications the characteristic impedance will be unimportant, as the line will not be correctly terminated, but at UHF and higher it is possible to use PCB tracks as microstrip transmission lines in properly terminated systems. If losses in such systems are to be minimized the PCB material must be chosen for low high frequency loss. This usually means the use of expensive teflon PCB material.

MICROSTRIP TRANSMISSION LINE




Figure 11.49

Where there is a break in the ground plane under a conductor the return current must flow around the break and both the

inductance and the vulnerability of the circuit to external fields are increased.

BREAKS IN GROUND PLANE RAISE INDUCTANCE




Figure 11.50

Where such a break is made to allow a crossover of two perpendicular conductors it would be far better if the second signal were carried across both the first and the ground plane by means of a piece of wire. The ground plane then acts as a shield between the two signal conductors, and the two ground return currents, flowing in opposite sides of the ground plane as a result of skin effects, do not interact.

With a multi-layer board both the crossover and the continuous ground plane can be accommodated without the need for a wire link. Multi-layer PCBs are expensive and harder to trouble-shoot than simple double-

sided boards but do offer even better shielding and signal routing. The principles involved remain unchanged but the range of layout options is increased.

Use of double-sided or multi-layer board with at least one continuous ground plane is undoubtedly one of the most successful approaches to the design of high performance mixed signal circuitry. Often the impedance of the ground plane is sufficiently low to permit the use of a single ground plane for both analog and digital parts of the system, but this does depend upon the resolution and bandwidth required and the amount of digital noise in the system.

SYSTEM GROUNDS

In systems where there are several PCBs grounding may be more of a problem. At first sight it would appear that the problem is similar to that of a single PCB where particular subsystems must be positioned so that large ground currents do not flow where ground noise must be minimized - in a multi-card system the grounds of individual PCBs must be interconnected so that such harmful interactions are minimized.

There are three problems with this. First of all there is far less opportunity for rearranging the physical layout of a system consisting of a few cards connected to a common backplane. Secondly many multi-card systems are designed to be reconfigured in a "mix 'n' match" arrangement to allow large numbers of system options - it can be impossible to predict what systems are going

to be required and to ensure that all of them are noise free. Finally, multichannel systems are likely to have higher ground currents than occur on single, relatively simple, PCBs - but these currents must flow in the higher impedances which are associated with the intercard connectors even when multiple ground pins are used.

The basic principles still apply: ground impedance must be as low as possible, high level and low level signals must be separated so that they do not interfere with each other, and capacitance and mutual inductance coupling must be avoided. Nevertheless, it must be accepted that situations can arise where it is not possible to transfer a high speed, high accuracy signal from one PCB to another without unacceptable signal degradation.

MULTIPLE CARD SYSTEMS

- Multiple card systems are likely to have higher ground currents and higher ground impedances than are found on a single PCB.
- It is therefore more difficult to transfer ground-referenced signals accurately between cards than across a PCB.
- In some cases it will be IMPOSSIBLE to transfer ground-referenced signals between PCBs without unacceptable loss of quality.

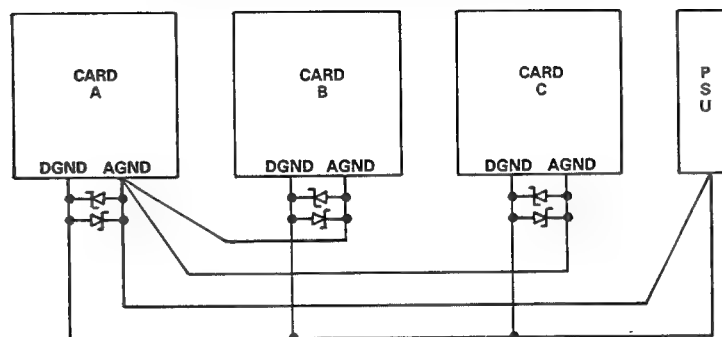
Figure 11.51

The best way of minimizing ground impedance in a multcard system is to use another PCB as a backplane and have a ground plane (or even two - one analog, one digital) on that mother card. If the earlier advice about multiple ground pins has been observed this arrangement is capable of excellent performance. Where there are several card cages (racks for PCBs) the ground planes of the several mother boards must be tied together and, probably, to the metal chassis holding the card cages - the

exact layout of the interconnections will depend on the overall system architecture.

If a mother board with a ground plane is not possible then the ground pins of the PCB sockets must be wired together, with due attention to probable current flows and common ground impedances, with heavy, multi-strand wire, having as low resistance as possible. In many cases the resulting ground screen will be tied to chassis ground at a number of points but it will sometimes be better to join them at a single star point.

STAR ANALOG GROUND IN A MULTICARD SYSTEM



- Schottky diodes protect cards in the event of loss of analog ground
- This grounding system may be inadequate at high resolution or where large ground currents flow
- This MAY permit accurate intercard transmission of ground referenced signals

Figure 11.52

It is not just the ground layout that is important in high performance mixed signal systems, the siting of different subsystems

and the routing of signals is most important in determining overall system performance.

SIGNAL ROUTING

It is evident that we can minimize noise by paying attention to the system layout and preventing different signals from interfering with each other. High level analog signals should be separated from low level analog signals, and both should be kept away from digital signals. We have seen elsewhere that

in waveform sampling and reconstruction systems the sampling clock (which is a digital signal) is as vulnerable to noise as any analog signal, but is as liable to cause noise as any digital signal, and so must be kept isolated from both analog and digital systems.

SIGNAL ROUTING IN MIXED SIGNAL SYSTEMS

- **Physically Separate Analog and Digital Signals**
- **Avoid Crossovers Between Analog and Digital Signals**
- **Be Careful with Sampling Clock and A/D Converter Analog Input Runs**
- **Be Careful with High Impedance Points**
- **Use Lots of Ground Plane**
- **Use Microstrip Techniques for Controlled Impedances**

Figure 11.53

If a ground plane is used, as it should in be most cases, it can act as a shield where sensitive signals cross. Figure 11.54 shows a good layout for a data acquisition system

where all sensitive areas are isolated from each other and signal paths are kept as short as possible. While real life is rarely as tidy as this the principle remains a valid one.

PCB FLOWCHART

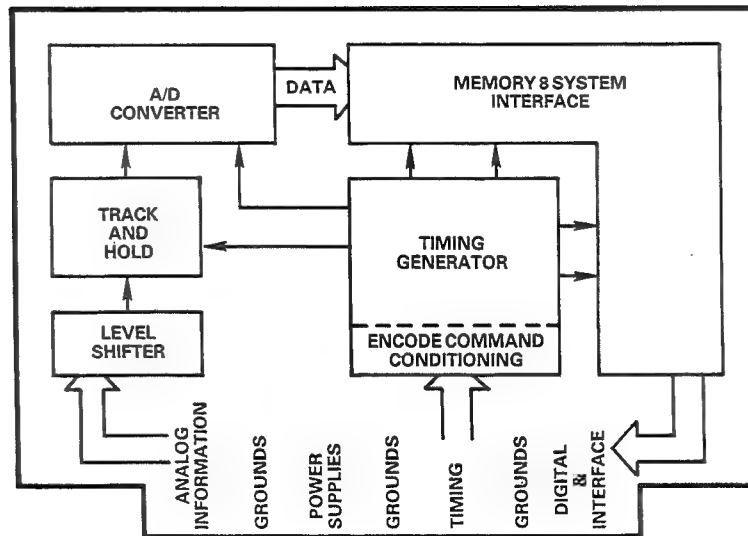


Figure 11.54

There are a number of important points to be considered when making signal and power connections. First of all a connector is one of the few places in the system where all signal

conductors must run parallel - it is therefore a good idea to separate them with ground pins to reduce coupling between them.

EDGE CONNECTIONS

- Separate sensitive signals by ground pins
- Keep down ground impedance with multiple (20-30% of total) ground pins
- Have several pins for each power line
- Critical signals may require a separate connector (possibly co-ax)

Figure 11.55

Multiple ground pins are important for another reason: they keep down the ground impedance at the junction between the board and the backplane. The contact resistance of a single pin of a PCB connector is quite low (of the order of 10 mohms) when the board is new - as the board gets older the contact resistance is likely to rise, and the board's performance may be compromised. It is therefore well worthwhile to afford extra PCB connector pins so that there are many ground connections (perhaps 20-30% of all the pins on the PCB connector should be

ground pins). For similar reasons there should be several pins for each power connection, although there is no need to have as many as there are ground pins.

Modern high performance mixed signal systems handle signals with resolutions of 8 bits at sampling rates of over 500 MHz and resolutions of 14 bits sampled at more than 10 MHz. Preserving signal integrity between cards in a multi-card system is extremely difficult at such performance levels and may be impossible.

DIFFERENTIAL TRANSMISSION MINIMIZES GROUND ERRORS

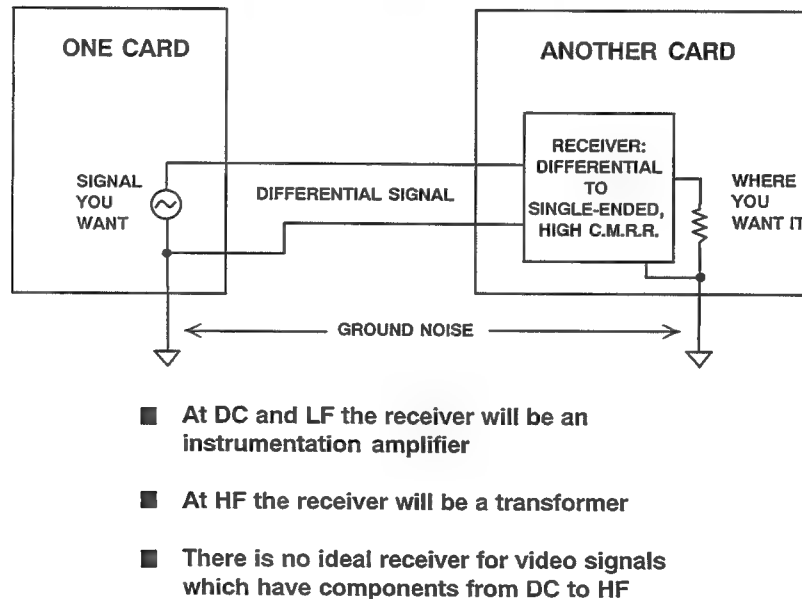


Figure 11.56

The use of balanced transmission lines can help but if the signal bandwidth extends to DC there will be a need for a very high

performance instrumentation amplifier at the receiving end to restore a ground referenced signal.

VIDEO SIGNAL TRANSMISSION

- It is often IMPOSSIBLE to transmit very broadband high accuracy signals between the PCBs of a multichip system without unacceptable loss of quality.
- In such cases the system must be reconfigured to allow all the analog processing to take place on a single PCB.
- It may be inconvenient, but it's the only way you'll get it to work!

Figure 11.57

The best, and in many cases the only, solution to problems of this sort, is to parti-

tion the system so that the highest quality signals are not transferred between boards.

POWER SUPPLIES

When we design an electronic circuit we generally assume that the power supplies provide noise-free power, at exactly the nominal voltage, with zero source impedance at all frequencies. This is rarely the case.

We also assume that the published power supply rejection figures (PSRR) for the devices which we use are valid at all frequencies from DC to light. This is rarely the case either.

POWER SUPPLY NOISE

- Long-term voltage variation
(Long-term variations in voltage or AC line voltage)
- AC Line noise
(Both 100/120 Hz ripple on rectifier output and transient noise on the AC line which passes to the DC output)
- Switching noise
(Digital noise from switching-mode power supplies)
- Power line noise transfer
(Unwanted signals which pass from one part of a circuit to another via the common power supply)

Figure 11.58

POWER SUPPLY NOISE

Every power supply is noisy. This noise may contain long-term voltage drift, line ripple at 100 or 120 Hz, high frequency spikes from switching regulators, or all of these at once. Power supplies also have finite output impedance, so that if a circuit draws a varying current the supply voltage will vary with the current - if two circuits are supplied from a common supply this provides a mechanism whereby one circuit may affect the other. Once we appreciate all these effects we can attempt to quantify them, and take steps to minimize their adverse effects on our systems.

Long-term supply voltage changes, whether due to battery voltage drop during life or line voltage variations, are rarely a problem since where such variations might cause difficulties the system will incorporate

a supply voltage regulator to keep variations within acceptable limits. Similarly ripple at twice the AC line frequency, and any spikes or HF noise which may enter the system via the AC supply, should not cause degradation of performance in a well-designed system: if the decoupling capacitors in the rectifier circuitry do not adequately minimize the effect, the series regulator almost certainly will. It is, however, always worthwhile to have a surge eliminator on the AC line input to any system - while such a circuit is unlikely to be needed in preventing normal line noise from corrupting system performance, it is essential to prevent occasional large surges (from lightning or similar causes) from doing actual damage to the power supply or the system that it is powering.

SWITCHING-MODE POWER SUPPLIES

The commonest type of power supply noise is switching noise. Switching power supplies are small, cheap, efficient and, in too many cases, extremely noisy! Not only do they generate conducted noise, they are also

efficient producers of capacitively coupled noise, magnetically coupled noise, and electromagnetically coupled noise. The best possible advice is not to use them.

SWITCHING-MODE POWER SUPPLIES

- **Generate every imaginable type of noise and some inconceivable ones as well!**
- **DO NOT USE THEM WHERE NOISE IS IMPORTANT**
- **If their use is unavoidable do not relax and enjoy it, *but* take extreme precautions against all forms of noise.**
- **Remember that a manufacturer's design change in a bought-in switching-mode power supply may alter its effects on your system noise without altering its published specification.**

Figure 11.59

It is, unfortunately, not always possible to avoid the use of switching power supplies. Where they must be used they must be treated with the gravest suspicion and all possible precautions should be taken to prevent their noise from corrupting the analog circuits that they power. Their input and output lines should be decoupled at all frequencies, they should be shielded to prevent external electric and magnetic fields from causing interference, and they should be sited as far as possible from sensitive circuits so that residual electric and magnetic fields are prevented by distance from doing serious damage.

Where switching supplies are used it is always worthwhile to remove them temporarily and supply the system with batteries or a low noise bench supply in order to determine if the system performance is being

compromised by the switching supply. It often is.

The noise transients on the output lines of switching supplies consist of voltage spikes of very short duration. As we have pointed out above, large capacitors, such as electrolytic or plastic film types, have quite considerable inductance and too high an impedance at HF to decouple such spikes satisfactorily. The best output filter for a switching supply will have high value capacitors to remove the low frequency noise which will also be present, and a pi filter using ceramic capacitors, with short leads having low impedance at HF, plus a series inductor (which may be a ferrite bead on the output line) to provide inductive blocking of the spikes. It is possible to buy such a pi filter as a single bulkhead mounted feedthrough component.

ELECTROMAGNETIC INTERFERENCE

RADIO FREQUENCY INTERFERENCE

Noise can enter a circuit as electromagnetic radiation. Circuits can also generate electromagnetic radiation which can interfere with electronic devices at quite considerable distances away. Recent legislation in the United States, the European Community and many other countries sets limits on the amount of interference generated and the vulnerability of circuits to such interference.⁷

This legislation, and the techniques needed to comply with it, are the subjects of many seminars and training courses, and an Analog Devices Application Note.⁸ It is not proposed to cover the topics in detail in this seminar.

ELECTROMAGNETIC NOISE GENERATION

- Circuits must be designed so that external E/M fields are minimized.
- This is done by shielding, decoupling, minimising the area of HF current loops and designing circuits which generate as little EMI as possible.
- IT'S NOT JUST A GOOD IDEA
- IT'S THE LAW!

Figure 11.60

However, the principles of minimizing external radiation are closely related to the principles of low noise design which we have already discussed: high frequency and high dV/dT signals should be screened with Faraday shields, the area of current loops should be minimized, conductors should be decoupled at HF wherever unnecessary HF signals might otherwise occur, and external wires should be isolated with inductors or ferrite beads.

It is still too common at seminars like this to encounter skepticism about the need to

protect circuitry from external electromagnetic fields. Even twenty years ago such skepticism was unjustified but today, when transmitters are ubiquitous, it is folly. Besides the more obvious broadcast, emergency and mobile radio services there are cellular and cordless telephones, radar, garage door openers and other remote controls, telemetry, and amateur and CB radio. For any designer to imagine that his circuit will never encounter a radio transmitter during its lifetime is folly on a grand scale.

ELECTROMAGNETIC NOISE INTERFERENCE

- The World is full of radio transmitters.
- Police, taxis, broadcast, amateur, CB, cellular and cordless telephones, telemetry and garage door openers.
- Do not imagine that your circuit will never encounter one.

Figure 11.61

This is particularly so because the design of circuits which are immune to electromagnetic radiation of reasonable levels is not particularly difficult. If every conductor

which leaves a PCB can be decoupled with a ceramic capacitor and a ferrite bead, it is probable that no further precaution is necessary.

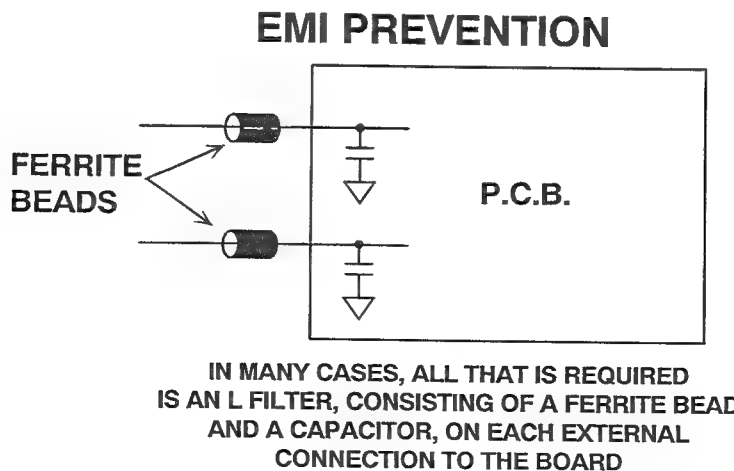


Figure 11.62

A few ports may be more vulnerable and require a pi filter rather than an L filter, and, of course, ports where an HF signal must actually enter or leave the board must be filtered to suppress other EMI but allow the signal to pass unaffected.

Boards which may be required to work in areas of high RF field should be screened with a conducting Faraday shield.

PHOTOELECTRIC EFFECTS

Light is also a form of electromagnetic radiation and can effect semiconductor devices. Every silicon P-N junction is a photodiode, although their efficiencies vary widely. Wherever devices are not screened from ambient light, photoelectric effects may be observed.


Nearly all integrated circuits are encapsulated in light-tight packages (EPROMs are an exception, and it is possible to measure threshold changes in EPROMs as light intensity is varied, but since they are digital devices, and remain in specification despite light level changes, the effect is unimportant).

Diodes, on the other hand, are frequently encapsulated in translucent glass packages. When illuminated by light from fluorescent

lamps, modulated at 120 or 100 Hz, they can act as a source of hum.

When the signal source of an op-amp contains an energizing voltage which is much higher than the op-amp supply it is common to use a diode and a current limiting resistor to protect the op-amp in the event of a sensor short-circuit. In normal operation the diode is reverse biased and contributes only its (low) leakage current to the circuit but should the sensor be short-circuited the resulting current will flow through the diode to the op-amp supply rather than destroy the op-amp. It is, of course, important to choose the resistor so that it neither degrades the noise performance of the system nor allows too much current to pass under fault conditions.

UNEXPECTED PHOTOELECTRIC EFFECTS IN SILICON JUNCTIONS CAN DEGRADE CIRCUIT PERFORMANCE



THE DIODE PROTECTS THE OP AMP UNDER FAULT CONDITIONS BY DIVERTING FAULT CURRENT (LIMITED BY R_p) TO THE SUPPLY RAIL. THE DIODE SHOULD NOT BE PHOTO-RESPONSIVE, OTHERWISE FLUORESCENT LIGHTING MAY MODULATE ITS LEAKAGE CURRENT AT 100/120Hz AND CAUSE HUM. USE A PLASTIC DIODE – NOT A GLASS ONE.

Figure 11.63

The European Applications Department of Analog Devices encountered such a system where about 10% of all the amplifiers built suffered from severe hum at twice the power line frequency. The customer, of course, blamed the op-amp for poor supply rejection but analysis showed that even when the circuit was powered from batteries the problem persisted. The cause eventually turned out to be fluorescent lighting affecting the protective diode - a 1N914 in a glass case.

About 10% of diodes from the particular manufacturer were quite active as photodiodes and when illuminated by fluorescent lights their leakage current was modulated at 100 Hz (this was a European problem) - and the 100 Hz was, of course, amplified with the sensor signal. Use of a black epoxy packaged diode provided a complete cure.

LOGIC

The majority of this section of our seminar has considered problems within the analog parts of mixed signal systems. Despite their much greater noise immunity, the digital parts of these systems can also suffer from designers' lack of consideration of basic laws. Common problem areas include bus interface issues, including fan-out and timing, for both

converters and DSP processors, the care and feeding of sampling clocks, and the design of systems which generate minimum noise (we have already discussed how to keep logic noise from affecting the analog parts of a system - this task becomes easier if the logic noise is minimized in the first place).

DIGITAL PROBLEM AREAS IN MIXED SIGNAL SYSTEMS

- Bus interface - fan-out
- Timing variations
- Sampling clock jitter
- Logic noise

Figure 11.64

FAN-OUT

All Analog Devices' DSP processors, and most of their DSP ADCs, have TTL-compatible CMOS logic ports. The inputs have $V_{il(max)}$ of 0.8 V and $V_{ih(min)}$ of 2.0 V, while the outputs have $V_{ol(max)}$ of 0.4 V and $V_{oh(min)}$ of 2.4 V at particular currents. The DSP processors are also rated for the capacitive load that they will drive without degradation of their timing.

In order to determine the fan-out of such devices it is necessary to consider the current that they are called upon to source and sink and the capacitance that they will see. This is done from the data sheets of the devices that they will be called upon to drive.

FACTORS LIMITING LOGIC FAN-OUT

- Maximum available source current (logic high):
Dominant factor for *resistive* loads
- Maximum available sink current (logic low):
Dominant factor for *TTL* loads
- Maximum permitted node capacitance:
Dominant factor for *CMOS* loads
- Node capacitance has contributions both from the input capacitances of gates on the node and from wiring and PC tracks associated with the node.

Figure 11.65

Consider a typical fast TTL gate, such as the 74F32 OR-gate. Its maximum input high current ($I_{ih(max)}$) is 20 μ A, its maximum input low current ($I_{il(max)}$) is 0.6 mA and its maxi-

mum input capacitance is 5 pF. An ADSP-2100 will source 1 mA when its output is high, it will sink 4 mA when its output is low, and it will drive capacitance of up to 100 pF.

ADSP-2100 DRIVE CAPABILITY

- Will Drive 100 pF
- Will drive 1 mA at logic 1 (≥ 2.4 V)
- Will sink 4 mA at logic 0 (≤ 0.4 V)

Therefore it will drive:

- 22 74ACT CMOS Gates (= 99 pF, ± 22 μ A)
- 10 74LS Schottky TTL Gates (= -4 mA)
- 7 74F Schottky TTL Gates (= -4.2 mA)
- 1 Grounded 2.4 K resistor (= 1 mA)
- Or any combination of loads which does not exceed a total capacitance of 100 pF, a total drain of 1 mA at logic 1, and a total source of 4 mA at logic 0. (Remember to allow for the capacitance of PCB tracks and wiring)

Figure 11.66

The ADSP-2100 will therefore drive the capacitance of 20 74F32 gates, it will drive the input current of fifty such gates in the logic 1 (high) state, but it will sink the input current of only 6.7 (in practical terms, 7) such gates. The lowest of these is evidently the fan-out which it will drive.

In typical systems it is likely that a device will be called upon to drive a mixture of devices, so the calculations will be more complex - but the basic principle will be the same. In most systems involving TTL the fan-out will be limited by the sink current, but in CMOS systems the node capacitance is likely to be the limiting factor. The above calculations do not consider the capacitance of the PC tracks and any cables which the device may be called upon to drive but such capacitance can sometimes be a limiting factor, and should always be considered, if only to be eliminated.

Most data converters have less powerful output stages than processors and their fan-out is lower. Additionally the return current of the output drive from a converter will flow in the system analog ground (for reasons discussed earlier in this section) and should therefore be kept as low as possible in order to minimize digital noise in the analog part of the system.

This is best achieved by using CMOS, rather than TTL logic. The DC input currents of CMOS are orders of magnitude lower than those of TTL. However, the input capacitances are comparable so the switching transients are not much lower. It is therefore advisable to buffer ADC outputs with an external buffer to minimize digital output currents from the ADC. Such a buffer will also help to isolate the ADC from digital noise in the rest of the system.

TIMING VARIATIONS

A common cause of malfunctions in digital systems, and particularly in the digital parts of mixed signal systems, is timing error, which often arises from failure to consider the effects of temperature variations on the system.

The specifications of converters, memories and processors all contain such parameters

as "set-up" and "hold" times. These are the times, respectively, that data must be present before a clock edge may occur, or that it must remain valid after the edge. At room temperatures many digital circuits are quite tolerant of operation with set-up and hold times which are shorter than the specified minimum - but at extremes of temperature they may be more demanding.

LOGIC TIMING VARIES WITH TEMPERATURE

- Specifications such as “set-up” & “hold” (the time a signal must be present before a strobe and the time that it must remain after one, respectively) can vary widely with temperature.
- A system designed with room temperature “typical” values may only perform properly at room temperature, if then.
- Designers **MUST** use min/max specifications at temperature extremes to ensure correct operation at all times.

Figure 11.67

Where a system consists only of digital circuitry it is likely (but not certain) that changes in input and output timing will behave similarly so that systems continue to function over temperature. Where ADCs or DACs are interfacing with digital systems the very different processes used for the converters may result in timing changes not tracking and performance, or even functionality, suffering.

Engineers designing mixed signal systems should always ascertain that the maximum and minimum timing specifications of all the circuits in their systems are compatible over the full temperature range of intended operation. Where there is any doubt buffers or monostables should be used as pulse extenders to ensure that all set-up and hold specifications are complied with.

SAMPLING CLOCK NOISE

As has been mentioned elsewhere in this seminar, phase noise on the clock of a sampled data system is indistinguishable from phase noise on the signal itself and it is therefore of critical importance to ensure that the sampling clock has sufficient spectral purity that its phase noise is less than the smallest component to be detected in the signal under analysis.

To achieve this the sampling clock should be isolated as much as possible from the

noise present in the digital parts of the system. In particular, buffers used for the sampling clock should, ideally, be on separate chips, with separately decoupled supplies, from the remainder of the digital system, and the sampling clock signal lines should not be sited where they can pick up digital noise from the rest of the system.

SAMPLING CLOCK NOISE

- Phase noise of the clock must be less than the minimum signal to be detected in the system.
- Therefore the sample clock signal must be protected from digital noise.

BUT

- Clocks are digital and can corrupt the analog part of the system.
- Therefore sampling clock lines must be kept separate from both the analog and the digital parts of the system.
- The sampling clock must use an oscillator with low phase noise.


Figure 11.68

Of course the sampling clock is itself a digital signal. It has as much potential for causing noise in the analog part of the system as any other digital signal. In fact, due to its presence in the converter and SHA sections of a system, it is generally the leading suspect for noise. We therefore see that a sampling clock is very inconvenient as it must be isolated from both the analog and digital parts of the system.


The sampling clock generator must also have adequate spectral purity. RC and other relaxation oscillators just will not do since amplitude noise in whatever circuit functions as a comparator will appear as phase noise on the output signal. LC oscillators have better phase noise, but the lowest noise is obtained with the use of a quartz crystal oscillator. For very high speed clocks a SAW (surface acoustic wave) oscillator is preferable.

A popular design of quartz crystal oscillator uses a resistor, one or more logic gates, a quartz crystal and a couple of capacitors. The design is not popular with engineers who understand quartz crystals or oscillators - such designs have bad phase noise and are liable to overdrive the quartz crystal (not enough to shatter it, as sometimes happened with self-excited crystal-controlled transmitters using vacuum tubes, but enough to affect its long-term stability). The only justification for the use of such oscillators is in watch and clock circuits where the low voltages involved minimize the overdrive and the phase noise is integrated over long periods and so is unimportant.

OSCILLATORS



Crystal oscillators built with logic gates have large phase noise.



Low noise crystal oscillators are easily designed with discrete components.

Figure 11.69

Ideally quartz oscillators should use discrete bipolar and FET devices in the circuits recommended by the crystal manufacturers. These circuits are optimized for

LOGIC NOISE

One of the most common causes of loss of performance in mixed signal systems is degradation of analog performance by noise from the digital parts of the system. We have already discussed at some length how this digital noise may be isolated from the sensitive analog parts of the system, but it is also worthwhile considering how this noise may be diminished at its source.

It is well-known that TTL is noisy. This is partly because the “totem pole” output stage structure acts as a short-circuit on the supply for a nanosecond or so during switching - giving rise to a large current spike, partly because the current flowing in the input changes, and changes quickly, between logic

0 and logic 1, and partly because the output swing, which takes place in a few nanoseconds, is several volts.

The output signal can then be amplified (possibly with a logic gate at this point) to drive the converters.

High speed CMOS does not have the change in input current (although there is a capacitance charging current pulse during switching, this is smaller) but may draw a supply current pulse during switching and certainly has a large output swing with a large dV/dT .

4000-Series CMOS is almost 20 years old and slow. It is also widely available, cheap, resistant to RFI, and quite remarkably noise free, since it has low output dV/dT and does not generate a supply current pulse.

LOGIC NOISE

- TTL has large voltage swings large, fast I/P & O/P current pulses and asymmetrical circuitry.
- HCMOS has large voltage swings large, fast O/P current pulses and symmetrical circuitry.
- 4000-Series CMOS is old, slow, cheap and very quiet.
- ECL has smaller voltage swings and smaller current surges than TTL & HCMOS even though it is faster.
- There is no single ideal logic family

Figure 11.70

ECL also draws almost constant current during switching (unless it is driving asymmetrical loads) and has much smaller output voltage swings than TTL or CMOS. Thus, although ECL is faster than TTL and CMOS, it tends to generate less noise.⁹

No single logic family is ideal for all applications (otherwise there would only be one logic family) but it is safe to conclude that TTL should not be used where its noise can corrupt precision analog circuitry but should be replaced by CMOS.

Where only low speeds are necessary 4000 CMOS has overwhelming noise advantages

PROBLEM AREAS

LIMITATIONS OF SPICE MODELLING

As we have seen, real electronic circuits contain many “components” which were not present in the circuit diagram but which are there because of the physical properties of conductors, circuit boards, IC packages, etc. These components are difficult, if not impossible, to incorporate into computer modelling software and yet they have substantial

but may not be available in all necessary configurations, and does not interface well with TTL (although it will interface with high speed CMOS families).

In high speed systems where noise is important ECL may offer noise advantages at the interface between the analog and digital parts of the system, even though high speed CMOS is capable of the speeds being used. It is not necessary to use ECL throughout the system - just where its lower noise is advantageous.

effects on circuit performance at high resolutions, or high frequencies, or both.

It is therefore inadvisable to use SPICE modelling or similar software to predict the ultimate performance of such high performance analog circuits. After modelling is complete the performance must be verified by experiment.

This is not to say that SPICE modelling is valueless - far from it. Most modern high performance analog circuits could never have been developed without the aid of SPICE and similar programs, but it must be remembered that such simulations are only as good as the models used and these models are not perfect. We have seen the effects of parasitic components arising from the conductors, insulators and components on the PCB, but it is also necessary to appreciate that the models used within SPICE simulations are not perfect models.

Consider an operational amplifier. It contains some 20-40 transistors, almost as many resistors, and a few capacitors. A complete SPICE model will contain all these components and probably a few of the more important parasitic capacitances and spurious diodes formed by the diffusions in the op-amp chip. This is the model that the designer will have used to evaluate the device during his design. In simulations such a model will behave very like the actual op-amp, but not exactly.

SPICE MODELLING

- **SPICE modelling is a powerful tool for predicting the performance of analog circuits.**

HOWEVER

- **Models omit real-life effects**
- **No model can simulate all the parasitic effects of discrete components and a PCB layout.**

THEREFORE

- **Prototypes must be built and proven before production.**

Figure 11.71

However, this model is not published, as it contains too much information which would be of use to other semiconductor companies who might wish to copy or improve on the design. It would also take far too long for a simulation of a system containing such models of a number of op-amps to reach a useful result. For these, and other, reasons

the SPICE models of analog circuits published by manufacturers or software companies are “macro” models, which simulate the major features of the component but lack some of the fine detail. Consequently SPICE modelling does not always reproduce the exact performance of a circuit and should always be verified experimentally.

SOCKETS

It is tempting to mount expensive ICs in sockets rather than soldering them in circuit - especially during circuit development.

Engineers would do well not to succumb to this temptation.

USE OF SOCKETS WITH HIGH PERFORMANCE ANALOG CIRCUITS

- **DON'T! (If at all possible)**
- **Use "Pin sockets" or "Cage jacks" such as Amp Part No: 5-330808-3 or 5-330808-6 (Capped & uncapped respectively).**
- **Always test the effect of sockets by comparing system performance with and without the use of sockets.**
- **Do not change the type of socket used without evaluating the effects of the change on performance.**

Figure 11.72

Sockets add resistance, inductance and capacitance to the circuit and may degrade performance to quite unacceptable levels. When this occurs, though, it is always the IC manufacturer who is blamed - not the use of a socket. Even low profile, low insertion force sockets cannot be relied upon not to degrade the performance of high performance (high speed or high precision or, worst of all, both) devices, and as the socket ages and the board suffers vibration the contact resistance of low insertion force sockets is very likely to rise. Where a socket must be used the least loss of performance is achieved by using individual pin sockets (sometimes called "cage jacks") to make up a multi-pin socket in the PCB itself.

It really is best not to use IC sockets with high performance analog and mixed signal

circuits. If their use can be avoided it should be. However at medium speeds and medium resolutions the trade-off between performance and convenience may fall on the side of convenience. It is very important, when sockets are used, to evaluate circuit performance with and without the socket chosen to ensure that the type of socket chosen really does have minimal effect on the way that the circuit behaves. The effects of a change of socket on the circuit should be evaluated as carefully as a change of IC would be and the drawings should be prepared so that the change procedures for a socket are as rigorous as for an IC - in order to prevent a purchase clerk who knows nothing of electronics from devastating the system performance in order to save five cents on a socket.

PROTOTYPING HIGH PERFORMANCE ANALOG CIRCUITRY

As we have seen, circuit board layout is part of the circuit design of all high performance analog circuits. Prototyping techniques derived from the “node” theory, while ideal for logic breadboarding at low and medium speeds, are quite unsuitable for any analog circuits, or even for very fast digital ones. Vector board and wire wrap prototyping will tell an engineer nothing about the behavior of a properly laid out version of the analog circuit.

The best technique for analog prototyping is to use a prototype of the final PCB - certainly no design is complete until the final PCB layout has been proved to give the required performance. Nevertheless this approach may be a little limiting where a number of different possibilities are to be evaluated, or for a multcard system.

PROTOTYPING MIXED SIGNAL CIRCUITRY

- **NEVER** use vector boards or wire-wrap for the analog parts of the system (they can be invaluable for data buses and address lines in the digital part).
- **Wherever possible avoid the use of sockets for analog ICs.**
- **Use a prototype of your final PCB layout as early as possible.**

Figure 11.73

In this case components should be mounted on a board having a continuous copper ground plane (ideally on both sides of the board, though while convenient this is not essential), with ground connections made to the plane and short point to point wiring made above and below it. The overall component placing and signal routing should be as close as possible to the planned final layout.

As we have already indicated, IC sockets can degrade the performance of analog ICs. While directly soldered components are ideal for prototyping, an IC socket made of pin sockets mounted in the ground plane board may be acceptable (clear the copper, on both sides of the board, for about 0.5 mm around each ungrounded pin socket - solder the grounded ones to ground on both sides of the board).

Allowing wiring to float in the air can be a little tricky. There is a breadboarding system which is conceptually very similar to that described above but which provides adhesive PC pads which stick to the ground plane and allow more rigid component mounting and wiring. This system is manufactured by Wainwright Instruments and is known as “Minimount” in Europe and “Solder Mounts” in the USA. The manufacturer’s and distributors’ addresses are given in the references at the end of this section.¹⁰

Manufacturer’s evaluation boards are also useful in system prototyping since they have already been optimized for best performance. Analog Devices offers many evaluation boards for a wide array of products. They offer the designer an excellent starting point for the layout.

When the prototype layout is transferred to a CAD system for PCB layout it is important to disable, or at any rate override where necessary, any automatic routing or component placing software. The criteria used by such software are more closely related to

“node” theory and aesthetically pleasing rows of components (which, admittedly, are also easier on automatic component placing machinery) than to optimizing stray inductance and capacitance and minimizing common ground impedances.

ADDITIONAL PROTOTYPING HINTS

- **Pay *equal* attention to signal routing, component placing and supply decoupling in *both* the prototype and the final design.**
- **Verify performance as well as functionality at each stage of the design.**
- **For “freehand” prototyping use a copper-clad board, mount components to it by their ground pins and wire the remaining connections point-to-point (use Wainwright Instruments’ Minimount/Solder Mount adhesive PC pads if aerial point-to-point wiring seems too fraught with peril).**

Figure 11.74

REFERENCES

1. Albert Einstein
2. "E.S.D. Prevention Manual" Available free from Analog Devices.
3. For a discussion of these effects see pp I-11 et seq. "Linear Design Seminar" notes. Analog Devices, October 1987.
4. "Electricity & Magnetism" by B.I. & B. Bleaney. OUP 1957, pp 23,24, & 52.
5. "An I.C. Amplifier User's Guide to Decoupling, Grounding and Making Things Go Right for a Change" by Paul Brokaw. Available free of charge from Analog Devices.
6. "Avoiding Ground Problems in High Speed Circuits", Jeff Barrow "R.F. Design" July 1989. AND "Grounding for Low- and High-Frequency Circuits", Paul Brokaw & Jeff Barrow, "Analog Dialogue" 23-3 1989. Free from Analog Devices.
7. International EMI Emission Regulations
Canada CSA C108.8-M1983 FDR VDE 0871/VDE 0875
Japan CISPR (VCCI)/PUB 22 USA FCC-15 Part J
8. "Design & Layout of a Video Graphics System for Reduced EMI" Bill Slattery & John Wynne (E1309-15-10/89). Free from Analog Devices.
9. "MECL System Design Handbook" 4th Edition, 1983 - Motorola
10. Wainwright Instruments Inc.
P.O.Box 2203, Capistrano Beach, CA 92624. Tel: 714-248-2490

Wainwright Instruments GmbH, Widdersberger Strasse 14,
DW-8138 Andechs-Frieding, Germany. Tel: +49-8152-2245

INDEX

INDEX

SUBJECTS

A

Abdominal imaging, ultrasound, II.32
 Absolute value circuit, II.4, 15-16
 op-amp, II.4
 AC linearity plot, III.25-26
 Active filter, III.37
 AD7528 Dual 8-Bit CMOS DAC Application Note, II.51
 Adams, R.W., VI.30
 Adaptive pulse code modulation system, See: ADPCM
 ADC, I.2, 6, II.20, 31-2, 34, 36, 39, III.3, 7-9, 25, 27, VII.1-2, 22
 3-bit,
 ideal, analog input, III.18
 non-ideal, analog input, III.18
 8-bit, subranging, IV.5
 12-bit, subranging, IV.6
 14-bit, recursive, IV.7
 AC linearity, III.26
 aperture delay time, III.27
 aperture jitter, III.27-28
 best-fit sinewave, III.15
 busy/interrupt, IX.4
 conversion complete, IX.4
 converter encoder, IV.2
 data acquisition, III.29
 data ready, IX.4
 data valid, IX.4
 DCS, IV.5-6
 demodulation, III.10-11
 digital output, III.17
 DMA, IX.4
 DMD, IX.4
 DNL, III.17
 DNL plot, III.25-26
 DSP applications, IV.1
 DSP interface, IX.14
 dual slope, IV.8

dynamic performance, III.21
 dynamic range, III.12
 dynamic specifications, III.21
 dynamic testing, III.14
 effective aperture delay time, III.27
 encode command, IX.4
 end-of-conversion, IX.4
 ENOBs, III.15, 21
 errors, III.18
 FFT output, III.14
 finite amplitude resolution, III.12
 flash, IV.3-4
 N-bit, IV.3
 performance, III.22
 FLB, III.23
 FPBW, III.23-24
 gain error, III.17
 ground, XI.37
 and harmonic distortion, III.17
 IMD, III.25
 INL, III.17
 input, ripple, IV.9
 integrator, IV.8
 charge/discharge, IV.8-9
 hysteresis and offset, IV.9
 leakage and offset, IV.9
 normal mode, IV.9-10
 sampling period, IV.9
 interfacing, IX.1
 intermodulation, III.24
 memory-mapped, IX.1
 missing codes, III.17
 offset error, III.17
 overvoltage recovery, III.30
 parallel, to DSP, IX.4
 parallel interface, IX.4
 peak spurious response, III.23
 quantization noise, III.5, 12
 read data, IX.4
 recursive subranging, IV.7

- resolution and dynamic range, III.12, 16
- S/N+D, III.15, 21
- sampling, III.21-22, IV.4
- sampling converter, IV.1
- SAR with SHA, IV.1-3
- serial, to DSP, IX.15
- serial interface timing, IX.18
- serial output, IX.17
- settling time, III.29
- SHA, III.24
- sigma-delta technology, IV.9
- sinewave, III.15
- SNR, III.12, 21
 - aperture jitter, III.28
- start-convert, IX.4
- state of the art, I.7
- static performance specifications, III.17
- static transfer characteristics, III.17
- subranging, IV.5
 - 8-bit, IV.5
 - 12-bit, IV.6
 - digital correction, IV.5
- successive approximation, IV.1-2
- testing, III.13
- THD, III.23
- theoretical RMS error, III.15-16
- transfer function, III.18
- transient response, III.29
- tri-state output, IX.26
- ADC/ADSP, IX.2
- ADC/ADSP-2100, parallel interface, IX.7
- ADC/ADSP-2101:
 - parallel interface, IX.4
 - DMACK, IX.6
 - parallel read interface, timing, IX.5
- ADC/TMS32020/C25, parallel interface, IX.7
- Adder:
 - instrumentation amp, II.2
 - op-amp, II.2
- ADDS-2101-SW DSP Software Development Tools, Data Sheet, VIII.19
- ADPCM, I.3
- ADSP-2100, drive capability, XI.53
- ADSP-2100 Family Applications Handbook, VIII.19
- ADSP-2100 Family Applications Handbook Vol. 4, X.30
- ADSP-2100 User's Manual/Architecture, IX.28
- ADSP-2100/2102 User's Manual/Architecture, IX.28
- ADSP-2100/ADSP-2100A Digital Signal Processor, Data Sheet, VIII.19
- ADSP-2100A, wait state generation, IX.29-32
- ADSP-2101:
 - with ADC, DAC, and DSP, IX.1
 - DMA, IX.4
 - DMS, IX.4
 - I/O port interface, IX.24
 - memory write timing, IX.9
 - memory-mapped peripherals, IX.26
 - output enable, IX.3
 - parallel interface, IX.2
 - parallel read timing, IX.3
 - parallel write timing, IX.10
 - in QAM, X.6
 - RD, IX.4
 - RD timing, IX.2
 - serial interface, IX.17, 22
 - serial port, IX.15
 - serial port interface, IX.26
 - serial port receive timing, IX.16
 - serial port transmit timing, IX.20
 - V.32 modem,
 - program memory, X.8
 - serial interface, X.9-10
 - wait state generation, IX.29-32
- ADSP-2101 Cross-Software Manual, VIII.19
- ADSP-2101 DSP Microcomputer, Data Sheet, VIII.19
- ADSP-2101 Emulator, Data Sheet, VIII.19
- ADSP-2101 Emulator Manual, VIII.19
- ADSP-2101 EZ-ICE Manual, VIII.19
- ADSP-2101 EZ-LAB Manual, VIII.19
- ADSP-2101 microcomputer, VIII.9
 - architecture, VIII.5-6
 - configuration, VIII.17

- DAG, VIII.6
- development system, VIII.18
- DMA bus, VIII.7
- DMD bus, VIII.7
- PMA bus, VIII.7
- PMD bus, VIII.7
- program sequencer, VIII.15
- result bus, VIII.7
- serial ports, VIII.16
- shifter, VIII.11
- single-instruction cycle, VIII.5
- system interface, VIII.17
- ADSP-2101 User's Manual, VIII.19
- ADSP-2101-EZ Tools, Data Sheet, VIII.19
- ADSP-2105 DSP Microcomputer, Data Sheet, VIII.19
- ADSP-2111 DSP Microcomputer with Host Port, Data Sheet, VIII.19
- ADSP-2111 User's Manual, VIII.19
- ADSP-21msp50:
 - GSM benchmarks, X.18
 - mixed signal processor, X.17
 - optimized DSP, X.16
 - specifications, X.17
- ADSP-21XX DSP Hardware Development Tools, Data Sheet, VIII.19
- ADSP-21XX DSP Software Development Tools, Data Sheet, VIII.19
- ADSP-28msp01:
 - integrated modem, analog front end, X.10
 - specifications, X.10-11
- ADSP-28msp02 CODEC serial port DSP interface, IX.26
- ADSP-28msp02 sigma-delta CODEC, IX.25
- ADSP-28msp02 sigma-delta voiceband CODEC, IX.25
- AGC loop, in disk drive read amplifier, II.5
- AGND, XI.36-38
- AIN, IV.7
- Aliasing:
 - dynamic range of sample, III.3-5
 - frequency domain effect, III.4
 - on dynamic range, III.4
 - time domain effect, III.3
- unwanted tones, III.5
- Allen, P.E., II.52
- ALU:
 - features, VIII.8-9
 - instructions, VIII.9
 - register, VIII.9
- Amplifier, signal conditioner, I.8
- Amplitude shift keying, See: ASK
- Analog beamforming, II.34
- Analog cellular radio:
 - FDMA, X.11-2
 - frequency reuse, X.11-12
 - frequency spectrum allocation, X.12
 - problems, X.11
 - TDMA, X.13
- Analog circuit:
 - ESD, XI.9
 - high frequency decoupling, XI.23
 - high frequency instability, XI.24
 - noise, XI.36-37
 - SPICE modeling, XI.58-59
- Analog computation, with multiplier, II.14
- Analog computer, II.21, 24
- Analog downconversion, III.10
- Analog filtering:
 - active, II.37-38
 - passive, II.37
- Analog ground, See: AGND
- Analog multiplexer, III.29
- Analog multiplier, II.8-9
 - logarithmic, II.9
- Analog signal:
 - aliasing, III.2-4
 - bandwidth, sampling, III.2
 - conditioning and processing, I.6
 - discrete sampling, III.2
 - moving average FIR, VII.5
 - routing, XI.44
- Analog signal processing, I.6, II.5
 - amplifier, II.1
 - filtering, II.37
 - mathematical operations, II.1
 - options, I.6
- Analog sound recording:

additive noise, X.21
degradation, X.21
nonlinearities, X.21
Analog-to-digital conversion, III.1
Analog-to-digital converter, See: ADC
Andreas, D., VI.30
Andrews, James R., III.43
Antialiasing filter, III.6, 42
 alpha value, II.44-45
 audio, II.49
 design, II.44, III.7
 FDNR, II.44-45, 47-49
 FDNR audio, II.49
 frequency value, II.44-45
 multiple feedback, II.44-46
 out-of-band attenuation, III.7
 oversampling, III.7
 requirements, III.6
 Sallen-Key, II.44-46
 See: Antialiasing filter
 selection, III.5-6
 simplification, III.8
 specifications, II.44
 state variable, II.44-45, 47-48
 tuning, II.44-45
Aperture delay time, III.27
Aperture jitter, III.27-28
Arithmetic logic unit, See: ALU
Arithmetic status register, See: ASTAT
ASK, modulation, X.3
ASTAT, VIII.14
Audio:
 DAC,
 glitch minimization, V.5
 segmentation, V.6
Audio filter, III.38
Audio line driver, I.15
 balanced, I.14
 differential, I.14
 high common-mode rejection, I.14
Audio line receiver, I.15
 differential, I.14
Audio power meter, multiplier, II.13
Automatic gain control, See: AGC

B

Balanced audio line driver, I.14
Balanced modulator, II.15-16
Baldwin, Eugene E., III.43
Barber, William L., II.51
Barrow, Jeff, XI.63
Basic diode log amp, II.21
Basic multiplier, II.8
Beamforming:
 analog, II.35
 digital, II.35
Bell, Barry A., III.43
Bennett, W.R., III.13, 43
BiMOS current switch, V.4
BiMOS DAC, V.4
Binary current scaling, DAC, V.1
Binary logic circuit:
 logic states, XI.3
 noise immunity, XI.3
Bingham, John, X.30
Blackman, R.B., III.43
Bleaney, B., XI.63
Bleaney, B.I., XI.63
Blinchikoff, H.J., II.51
Blood, signal attenuation, II.34
BMS, DSP, parallel interface, IX.1
Boot memory select, See: BMS
Boser, B., VI.30
Boyd, I., X.30
Broadband Amplifier Applications, II.51
Brodersen, Gray, III.44
Brokaw, Paul, XI.63
Brown, Edmund R., II.51
Budak, A., II.51
Buffer latch, as Faraday shield, XI.21
Burton, Phil, V.19
Butterfly calculation, VII.30
Butterworth filter, II.39, VII.19
 design, II.44
 maximally flat filter, II.39

C

- Cable driver:
 - inverting, I.16
 - noninverting, I.16
- Calhoun, George, X.30
- Capacitance, XI.15
 - stray, XI.15-16
 - tuned circuit, XI.31
- Capacitive coupling equivalent circuit, XI.16
- Capacitive noise, XI.16
- Capacitive shield, XI.17
- Capacitor:
 - dielectric absorption, XI.24
 - electrolytic, leakage, XI.22
 - equivalent circuits, XI.21
 - inductance, XI.22
 - leakage, XI.22
 - parasitic effects, XI.21
 - series/loss resistance, XI.22
 - SHA,
 - dielectric absorption, XI.25
 - leakage, XI.22
- Cascaded gain stage, II.31-32
- Cascaded limiting amplifier, II.25
 - gain, II.26
 - response curve, II.26
- Cascaded semi-limiting amplifier, II.21, 25
- Caves, J.T., III.44
- CB processes, DAC, V.1-2
- CCITT, recommendations, X.3
- CD player:
 - 18-bit,
 - 8X oversampling,
 - 3-pole antialiasing filter, X.27-28
 - reconstruction electronics, X.26
 - 20-bit,
 - 8X oversampling,
 - reconstruction electronics, X.27
 - THD+N, X.28-28
- ADC, X.22
 - antialiasing filter, X.25
 - audio DAC, CD, X.28-29
 - design,
 - digital filtering effects, X.26
 - oversampling effects, X.26
 - first generation, reconstruction electronics, X.24
 - oversampling, SNR, X.25
 - read electronics, X.23-24
 - second generation,
 - 18-bit, 8X oversampling, X.25
 - oversampling, X.25
 - serial data, X.23
- CDMA:
 - DSP, signal transmission, X.13
 - military, spread spectrum, X.13
- Cellular telephone, X.11
- Ceramic IC:
 - capacitive effects, XI.18
 - Faraday shield, XI.17
- Charpentier, A., VI.30
- Chebyshev filter, II.39, VII.19
 - comparison with FIR, VII.10
- Circuit noise, XI.36-37
- CISC:
 - architecture, VIII.1-2
 - programming, VIII.1
- Clarke, Robert M., II.51
- CMOS:
 - logic fan-out, XI.53-54
 - noise, XI.57-58
- CMOS DAC, V.3
 - current-steering, V.3
 - voltage-switching, V.3
- CMOS DAC Application Guide, II.51
- CMOS switch, III.37
- CMRR, I.10
- CNI, in receiver, X.16
- Code division multiple access, See: CDMA
- CODEC, speech, X.14
- Coleman, Brendan, III.43
- Colotti, James J., III.43
- Comfort noise insertion, See: CNI
- Common-mode rejection ratio, See: CMRR
- Compact disk, See: CD
- Compact disk player, V.1
 - sigma-delta DAC, VI.22

Complementary bipolar, See: CB
Complex-instruction-set computer, See: CISC
Computation, log and antilog circuits, II.9
Conductor, capacitive noise, XI.16
Continuous grounded conductor, XI.39
Copper foil, XI.3
Cosier, G., X.30
Counts, Lew, II.51
Current feedback amplifier, II.15, 17
Curtin, Mike, X.30

D

DAC, I.2, II.48, III.7, 9, IV.7, VII.1

3-bit,

ideal, digital input, III.20

non-ideal, digital input, III.20

6-bit, V.17

12-bit, architecture, V.15-16

12-bit multiplying, V.16

14-bit,

deglitching, V.14

DSP interface, V.8

functional diagram, V.7

segmentation, V.7

specifications, V.8

16-bit,

glitch impulse, V.10

monotonicity, V.9

MSB, V.9

segmentation, V.8-9

segmented architecture, V.9

specifications, V.10

18-bit, V.6

oversampling, V.6

SNR, V.6

specifications, V.6

18-bit audio, segmentation, V.6

20-bit,

digital offset, V.11

input frequency, V.12

oversampling, V.12

segmentation, V.11

THD+N, V.12

20-bit audio,

digital offset, V.11

input frequency, V.12

segmentation, V.11

THD+N, V.12

architecture, V.1

audio, V.5-6

ac parameters, X.28

clock input, X.28

for, CD, X.28-29

sampling rates, X.28

serial input, X.28

BiMOS, V.4, 8

binary division, V.1

bipolar, V.2

CB processes, V.1-2

CMOS, V.3-4, 7

architecture, V.15-16

FET switches, V.4

construction, V.1-2

current switching, V.1-2

current-steering CMOS, V.3

deglitcher, III.35-36, V.13-14, VI.22

digital filtering, truncation noise, X.28

digital offset, V.10

DNL, III.19

errors, III.19-20

for DSP, V.1

DSP interface, IX.14

DSP parallel interface, IX.27

dynamic performance, III.30-31

dynamic range, III.12

dynamic specifications, III.30

ECL bit switch, V.2

evaluation, V.1

FFT testing, III.34

finite amplitude resolution, III.12

gain errors, III.19

glitch, III.34

glitch impulse, III.33

glitch impulse area, III.32-33

glitch minimization, V.5

glitch reduction, V.10

ground, XI.37

- harmonic distortion, III.33
- high performance applications, V.1
- INL, III.19
- interfacing, IX.1
 - with DSP chip, V.1
- linear, V.17
- minimum glitch, V.5
- MSB, V.11
 - output, IV.1
- multiplying, See: MDAC
- multiplying with digital input, See: LOGDAC
- non-monotonicity, III.19-20
- offset errors, III.19
- output, IV.7
- oversampling ratio, V.1
- parallel, DSP interface, IX.10-11
- PNP differential pairs, V.1
 - as potentiometer, V.15
 - as programmable gain amplifier, V.15
- raster scan display, V.1
- reference voltage, V.15
- resolution and dynamic range, III.12
- sample-and-hold, V.13
- serial, DSP interface, IX.19
- settling time, III.31-32, IV.3
 - refined, III.32
- SHA, V.13
- SHA deglitcher, III.36, V.13
- sin X/X rolloff, III.36
- SNR, V.1
- static performance specifications, III.19
- static transfer characteristics, III.19
- static transfer function, III.19-20
- switching glitch, V.13
- THD, III.35, V.1, 11-12
- track-and-hold, V.13
- TTL bit switches, V.2
- voiceband, V.1
- voltage switches, V.1
- voltage-switching CMOS, V.3-4
- DAC/ADSP-2100, parallel interface, IX.13
- DAC/ADSP-2101:
 - parallel interface, IX.8, 11-12
 - parallel write timing, IX.12
- DAC/TMS32020/C25, parallel interface, IX.13
- DAG, VIII.12-13
 - address pointer, VIII.6
 - addressing, VIII.14
 - features, VIII.13
 - FFT, VIII.12
 - modulus logic, VIII.13
- DAT, ADC, specifications, X.22
- Data acquisition, I.3
 - using ADC, III.29
- Data address generator, See: DAG
- Data bus, high speed, digital noise, XI.20
- Data memory acknowledge, See: DMACK
- Data memory address, See: DMA
- Data memory data, See: DMD
- Data memory select, See: DMS
- Data separator, II.6
- Dattorro, J., VI.30
- Davis, Gary, X.30
- DC converter, II.18
- DCS, IV.6
- DDS, V.1
- Decimation, III.8
- Deglitcher, III.33, V.14
- Del Signore, B.P., VI.30
- Demodulation:
 - analog, III.10
 - direct IF to digital, III.11
- Detector, fast peak, II.4
- DFT, VII.22-24
 - 2-point, FFT decimation, VII.24
 - 8-point, FFT decimation, VII.25
 - butterfly calculation, VII.24
 - equation, VII.23
 - relation to FFT, VII.31
 - signal analysis, VII.24
- DGND, XI.36-38
- Dielectric absorption, capacitor, XI.24
- Differential amplifier, I.10, III.40
- Differential integrator:
 - active, III.41
 - SCF equivalent, III.41

- Differential non-linearity, See: DNL
- Differential receiver input, I.10
- Differentiator, II.2
 - op-amp, II.2
- Digital attenuator, V.18
- Digital audio, I.5, V.1
 - DAC, V.11
- Digital audio oversampling filter, II.49
- Digital audio studio recording, X.20
 - channel processing, X.20
 - D-range distortion, X.22
 - DAT players, X.22
 - DSP signal conversion, X.21
 - FIR filter, X.21
 - mixing, X.20-22
 - multiplier, X.21
 - performance specifications, X.22
 - sampling ADC, X.23
 - sigma-delta audio ADC, X.23
 - SNR, X.22
 - standards, X.22
 - system, X.21
 - techniques, X.21
 - THD+N, X.22
- Digital audio studio system, I.5
- Digital audio tape, See: DAT
- Digital beamforming, II.34
- Digital circuit, noise, XI.36-37
- Digital correction, in subranging ADC, IV.5
- Digital downconversion, III.11
- Digital filtering, VII.1-2
 - ADC, VII.3
 - advantages, VII.2
 - FIR filter, VII.2-3
 - limitations, VII.3
 - and shaped quantization noise, VI.9
 - video, VII.3
- Digital ground, See: DGND
- Digital mobile radio, V.1, X.11
 - CDMA, X.13
 - DSP, X.13
 - TDMA, X.13
- Digital noise, high speed data bus, XI.20
- Digital offset, glitch reduction, V.10
- Digital signal, routing, XI.44
- Digital signal processing, See: DSP
- Digital Signal Processing Applications Using the ADSP-2100 Family, VIII.19
- Digital system, timing variation, XI.54
- Digital voltmeter, IV.8
- Digital-to-analog converter, See: DAC
- Digitally controlled AGC system, V.18
- Digitally corrected subranging, See: DCS
- Diode-connected transistor, II.21, 24
- Direct digital synthesis, See: DDS
- Direct IF to digital demodulation, III.11
- Discontinuous transmission, See: DTX
- Discrete Fourier Transform, See: DFT
- Discrete time sampling, III.1
- Disk drive read, II.5, 7
 - amplifier, II.5-6
- Divider, from multiplier, II.14
- DMA:
 - ADC, IX.4
 - ADSP-2101, IX.4
- DMA bus, ADSP-2101 microcomputer, VIII.7
- DMACK, ADC/ADSP-2101, parallel interface, IX.6
- DMD, ADC, IX.4
- DMD bus:
 - ADSP-2101 microcomputer, VIII.7
 - MAC, VIII.9
- DMS:
 - ADSP-2101, IX.4
 - DSP, IX.1, 8
 - DSP parallel interface, IX.1
- DNL plot, III.25-26
- Doernberg, Joey, III.43
- Downconversion:
 - analog, III.10
 - direct IF to digital, III.11
- DSP56000, serial interface, IX.18, 23
- DSP, I.2, 6, II.34-35, III.8, 33, 37, VII.12, 14-16, 19, 22-23, 30
 - ADC interface, IX.1
 - ADC serial interface, IX.15
 - algorithm, VIII.1-2
 - aliasing, III.9

ALU, VIII.3
 applications, III.16-17
 applications and dynamic range, III.16
 architecture, VIII.2
 bit-reversing, VII.24
 chip select, IX.8
 circular buffer, VII.8, VIII.4
 clock frequencies, IX.26
 code optimization, VIII.1
 CODEC interface, IX.23
 DAC interface, IX.1
 decimation, III.8
 demodulation, III.10
 digital filtering, VII.1
 DMS, IX.1, 8
 dual operand fetch, VIII.3-4
 dynamic analog signal, III.30
 dynamic range, III.16-17
 extended dynamic range, VIII.3
 fast arithmetic, VIII.3
 fast MAC, VIII.3
 FIR filter, VIII.4
 glue logic, IX.1, 16, 27
 hardware, VIII.1
 Harvard architecture, VIII.4
 I/O port interface, IX.23-24
 interfacing, IV.3
 kernel equation, VIII.2
 limitations, VII.1
 MAC, VIII.2
 memory address bus, IX.1
 methods, I.6
 MSB, IX.15
 options, I.6
 oversampling, III.8
 parallel interface, IX.1, 27
 BMS, IX.1
 DMS, IX.1
 IRQ, IX.1
 PMS, IX.1
 writing to DAC, IX.8
 parallel peripheral device write interface,
 IX.9
 processor, VIII.2

RAM, VIII.4
 RD, IX.1
 read process, IX.1
 reading data, IX.1
 RFS, IX.15
 sampled data, III.1
 signal conditioning, III.1
 serial interface, IX.14, 27
 serial port, IX.27
 serial versus parallel interface, IX.26
 in TDMA and CDMA, X.13
 time sampling, analog signals, III.2
 undersampling, III.9
 WR, IX.8
 zero overhead looping, VIII.3
 DSP interface:
 parallel DAC, IX.10
 serial DAC, IX.19
 DSP/ADC, designs, III.21
 DSP/DAC:
 functions, IX.20
 signal requirements, IX.21
 timing relationship, IX.21
 DTX, X.15
 CNI, X.16
 functions, X.16
 SID, X.16
 VAD, X.16

E

E.S.D. Prevention Manual, XI.63
 Eckbauer, F., VI.30
 ECL, noise, XI.58
 Effective aperture delay time, III.27
 Effective number of bits, See: ENOBs
 Eidi, Fares, VII.32
 Einstein, Albert, XI.63
 Electromagnetic force, See: EMF
 Electromagnetic interference, See: EMI
 Electromechanical piezoelectric transducer,
 II.32
 Electrostatic damage, See: ESD
 Elliptical filter, VII.19

- degraded phase response, II.39
 - passband ripple, II.39
 - stopband ripple, II.39
 - EMF, ground loop, XI.35
 - EMI, XI.49-50
 - noise generation, XI.49
 - photoelectric, XI.51
 - in silicon junctions, XI.51
 - prevention, XI.50
 - radiofrequency, XI.49
 - ENCODE command, IV.7
 - Encoder-decoder, II.7
 - Engelhardt, E., VI.30
 - ENOBs, III.12, 21-22, 24
 - ADC, III.15, 21
 - calculation, III.13, 15
 - sinewave curvefit, III.15-16
 - SNR, III.13
 - Error, aliasing, III.40
 - Error correction, Viterbi decoding, X.9
 - ESD:
 - in circuitry, XI.8
 - dust, XI.10
 - low temperature, XI.9
 - protection, XI.8-10
 - Explicit RMS computation, II.19
- F**
- Faraday shield, XI.16-18
 - buffer latch, XI.21
 - Faraday's Laws, XI.2
 - Fast Fourier Transform, See: FFT
 - Fast peak detector, II.4
 - FDM data, III.9
 - FDM signal, III.10
 - super-Nyquist sampling, III.10
 - FDMA, analog cellular radio, X.11-12
 - FDNR, filter configuration, II.41, 43
 - Ferguson, P.F., Jr., VI.30
 - FFT, III.11, 13-14, 33, VII.22-23
 - algorithm, VII.31
 - benchmark processing time, VII.28
 - block floating point, VII.30
 - computational speed, VII.31
 - DAG, VIII.12
 - data scaling, VII.30
 - decimation, DFT, VII.24-25
 - design, VII.27
 - DSP, hardware, VII.27
 - endpoint discontinuities, VII.31
 - as fast implementation of DFT, VII.31
 - hardware, VII.26
 - N-point, computation, VII.26
 - output, ADC, III.14
 - plot, IMD, III.25
 - real-time processing, VII.31
 - real-time speech analysis, VII.27
 - relation to DFT, VII.31
 - resolution, VII.31
 - signal analysis, VII.22
 - sinewave,
 - integral number of cycles, VII.28
 - non-integral number of cycles, VII.29
 - spectral analysis, VII.27, 29
 - spectral leakage, VII.29
 - testing, DAC, III.34
 - windowing, VII.29
- Filter:
- active, II.37-38, III.37
 - configurations, II.43
 - active lowpass RC, III.41
 - SCF equivalent, III.41
 - allpass, II.37
 - analog, oversampling, III.9
 - analog lowpass, Nyquist sampling, III.8
 - bandpass, II.37
 - bandstop, II.37
 - Butterworth, II.39
 - capacitance, III.40
 - Chebyshev, II.39
 - class, II.37
 - cutoff frequency, II.38-39
 - design, II.38
 - digital, oversampling, III.9
 - elliptical, II.39
 - FDNR, II.43, 47
 - first-order continuous-time active lowpass,

- III.41
 - highpass, II.37
 - lowpass, II.37
 - conversion to highpass, II.41
 - minimum passband attenuation, II.38-39
 - multiple feedback, II.40-43, 46
 - normalized passive, II.45
 - order, II.38-40
 - parameters, II.38
 - passband ripple, II.38-39
 - passive, II.45
 - normalized values, II.45
 - poles, II.40, III.5, 7
 - programmable state variable, II.48
 - RC, III.41
 - ripple, II.39, 49
 - Sallen-Key configuration, II.40, 43, 46, 50
 - SC, III.40-41
 - sections, II.40
 - seven-pole FDNR 20kHz antialiasing, II.49
 - specifications, II.39-40
 - state variable configuration, II.40-43, 47-48
 - stopband frequency, II.38-39
 - switched capacitor, III.37-40
 - transformation, II.41
 - tuning table, II.44
 - variable bandwidths, III.40
 - wideband Sallen-Key, II.50
- Filtering, II.18
 - digital, See: Digital filtering techniques, III.37
- Finite amplitude resolution:
 - by ADC or DAC, III.12
 - from quantization, III.1
 - LSB, III.12
- Finite impulse filter, See: FIR
- FIR filter, III.8
 - 4-tap, VII.9
 - 69-tap, processor time, VII.15
 - 91-tap, VII.10
 - passband ripple, VII.11-23
 - performance, VII.11
 - audio,
 - CAD design, VII.13-17
 - coefficients, VII.14
 - filter type, VII.14
 - frequency, VII.14
 - gain, VII.14
 - quantization bits, VII.15
 - taps, VII.15
 - audio lowpass, VII.13
 - coefficients, VII.8
 - compared to IIR filter, VII.22
 - data memory addressing, VII.9
 - decimation, VII.19
 - design, VII.6, 8
 - CAD techniques, VII.11-12
 - input, VII.12
 - output, VII.12
 - passband, VII.12
 - stopband attenuation, VII.12
 - theorem, VII.10
 - wordlength, VII.12
 - digital form, VII.3-4
 - direct form, VII.13
 - DSP, VIII.4
 - circular buffering, VII.8
 - equations, VII.8
 - Equiripple FIR design, VII.13
 - filter coefficients, VII.12
 - frequency domain, VII.6-7
 - frequency response, VII.12, 16
 - glitch, VII.22
 - impulse response, VII.7, 10, 12, 17
 - linear phase, VII.18
 - highpass filter, VII.18
 - lowpass filter, VII.18
 - taps, VII.18
 - moving average, VII.4
 - analog signal, VII.5
 - convolution, VII.3, 5
 - sin X/X response, VII.3, 6
 - properties, VII.19
 - step response, VII.17
 - structure, VII.9
 - symmetrical coefficients,
 - highpass filter, VII.18

lowpass filter, VII.18
taps, VII.6, 12
time domain, VII.6-7
transfer function, VII.6
Fisher, J., VI.30
Flash ADC, IV.3-4
 performance, III.22
Flash converter, II.31, 36, IV.5
 for HDTV, IV.4
 sampling ADC, IV.4
FLB, of ADC, III.23-24
Four quadrant multiplier, II.11
Fourier transform, VII.7
 Z-transform generalization, VII.7
FPBW, III.21, 24
 of ADC, III.23
Freeman, D.K., X.30
Frequency dependent negative resistance,
See: FDNR
Frequency division multiple access, See:
FDMA
Frequency division multiplexed, See: FDM
Frequency shift keying, See: FSK
FSK, modulation, X.3
Full wave rectifier, II.4
Full-linear bandwidth, See: FLB
Full-power bandwidth, See: FPBW
Fullscale converter, bit sizes, XI.3
Function generator, for inverse function,
II.14
FWR, II.5-6

G

Galand, C., X.30
Ganesan, A., VI.30
Gaussian minimum shift keying, See: GMSK
General impedance converter, See: GIC
General switched telephone network, See:
GSTN
Ghausi, M.S., II.52, III.44
GIC, filter configuration, II.41, 43
Gilbert, Barrie, II.51
Gilbert Cell, II.10-11

Glitch energy, III.32
Glitch impulse, III.33-34
 area, III.32-33
Glitch reduction:
 by segmentation, V.5
 techniques, V.15
GMSK:
 I/Q RF demodulator, X.19
 spectral leakage minimizing, X.18
Gold, Bernard, III.44, VII.32
Gray, G.A., III.43
Ground, current noise, XI.34
Ground loop, XI.34-35
 and EMF, XI.35
Ground noise, XI.34
Ground plane, XI.40
 breaks, inductance, XI.41
 mixed signal system, XI.44
 in PCB, XI.39
 residual resistance, XI.39
Grounding, XI.32
 ideal, XI.33
 realistic, XI.33
GSM, X.11, 13-14
 baseband I/O port, X.18-20
 specifications, X.20
 cylinders, X.15
DAC, GMSK-coded ROM, X.19
discontinuous transmission, X.14
downconversion, X.18
DTX, X.15
equalization, X.14
FIR filter, X.18-19
frequency-time allocations, X.19
modulation, GMSK, X.18
modulation/demodulation, X.14
sigma-delta ADC, X.18-19
speech compression, X.15
speech encoder and decoder, X.14
TDMA, X.18
transcoder, X.14
upconversion, X.18
Viterbi decoder, X.14
GSTN, V.32 modem, X.4

Guard ring, insulation resistance, XI.7

H

Hanning window, VII.29-30
 Harmonic distortion, III.21, 33
 and ADC performance, III.17
 Harris, Frederic J., III.43, VII.32
 Harris, Steven, VI.30
 Harvard architecture, DSP, VIII.4
 Haspeslagh, J., X.30
 HDTV, I.3
 flash converter, IV.4
 Heise, B., VI.30
 Hellwig, K., X.30
 Higgins, Richard J., I.18, III.44, VII.32
 High frequency decoupling, analog circuit, XI.23
 High frequency instability, analog circuit, XI.24
 High performance analog circuit, prototyping, XI.61
 High speed cable driving, configurations, I.16
 High Speed Design Seminar (1990), I.18, II.51, III.44, V.19, VII.32
 High-definition television, See: HDTV
 High-gain limiting amplifier, II.15-16
 Higher order modulator loop:
 dynamic range, VI.12
 idling pattern, VI.12
 stability, VI.12
 Hodges, David A., III.43-44
 Hofmann, R., X.30
 HP Journal, III.43
 HP Product Note, III.43
 Huelsman, L.P., II.52
 Hughes, Richard Smith, II.51

I

ICNTL, VIII.14
 IEEE Trial-Use Standard for Digitizing Waveform Recorders, III.44
 IFC, VIII.14
 IIR filter, VII.19

analog counterpart, VII.19
 biquad, VII.19
 compared to FIR filter, VII.22
 digital, VII.19
 efficiency, VII.21
 feedback, VII.21
 first-order, lowpass, VII.20
 general equation, VII.19, 21
 glitch, VII.22
 minimum sampling period, VII.21
 non-linear phase, VII.21
 properties, VII.21
 second order, VII.19
 second-order, lowpass, VII.20
 throughput, VII.21
 IMASK, VIII.14
 IMD, III.25
 FFT plot, III.25
 Implicit RMS computation, II.20
 Inductance:
 ground plane breaks, XI.41
 mutual, XI.26
 signal routing, XI.26
 stray, XI.25
 wire, XI.25
 Inductive coupling:
 principles, XI.27
 ringing, XI.29
 Inductor:
 parasitic effects, XI.30
 Q, XI.32
 saturation, XI.30-31
 Industrial process control, sigma-delta ADC, VI.16, 19
 Infinite impulse response, See: IIR
 INL:
 ADC, III.17
 DAC, III.19
 Input signal, oversampling, III.8
 Instrumentation amp, I.10-11
 3 op-amp, I.11
 adder, II.1
 common mode voltage, I.11
 high performance, I.12

high performance-high speed, I.12
high speed, I.12-13
pulse response and settling time, I.13
pulse response, I.13
settling time, I.13
subtractor, II.2
wideband, I.12-13
Insulator, leakage, XI.6
Integral non-linearity, See: INL
Integrated circuit, III.37, 39
ESD, XI.9
Integrated circuit log amp, II.28-29
cascaded, II.30
dc coupling, II.29
features, II.29
laser-trimmed, II.28
transfer function, II.29
Integrator, op-amp, II.3
Intercept voltage, II.20
Intermodulation distortion, See: IMD
International EMI Emission Regulations, XI.63
International Telegraph and Telephone Consultative Committee, See: CCITT
Interrupt control register, See: ICNTL
Interrupt force and clear, See: IFC
Interrupt mask register, See: IMASK
Inverting adder, II.1
Inverting op-amp, II.21, 23
IRQ, DSP parallel interface, IX.1
Isolation amplifier, I.17
three-port, I.17

J

Johnson noise, resistor, XI.14-15
Jones, Ralph, X.30

K

Karagozyan, Kapriel, IX.28
Kaufman, Sid, III.43
Kelvin feedback, XI.6
Kester, W.A., III.44
Kirchoff's Law, XI.2, 26, 32

Kitchin, Charles, II.51
Koch, R., VI.30

L

Laker, K.R., II.52, III.44
Laplace transform, VII.7
Laser trimming of resistor, II.11
Leakage, insulator, XI.6
Leakage resistance, PCB, XI.7
Least significant bit, See: LSB
Lee, Hae-Seung, III.43
Lee, Wai Laing, VI.30
Lenz's Law, XI.2
Line driver, balanced, I.13-15
Linear coded ADC and DAC, VI.25
Linear predictive coding, See: LPC
Log amp, II.9, 20-24, 28, 31-32
basic, II.22
broadband operation, II.30
cascaded, II.30-32
detector, II.22, 27
diode/op-amp, II.23
disadvantages, II.10
error curve, II.29
with flash converter, II.31
full-wave detector, II.21
gain, II.26
intercept voltage, II.20
key parameters, II.27
logarithmic transfer, II.29
multi-stage, II.25-26
specifications, II.26-27
successive detection, II.27-28
transfer, II.21
transistor, II.24
true, II.23, 32
Log strip, II.28
Log video amp, II.23, 31-32
Logarithmic amplifier, See: Log amp
Logarithmic converter, II.20
LOGDAC, I.6, V.17
8-bit, V.18
applications, V.18

LSB, V.17
 ROM decoder, V.17
 Logic:
 fan-out, limitations, XI.52-53
 mixed signal systems, XI.52
 noise, XI.57-58
 sampling clock noise, XI.55-56
 timing,
 temperature variation, XI.55
 variation, XI.54
 Low frequency measurements, sigma-delta
 ADC, VI.16, 19
 LPC, cylinders, X.15
 LPC algorithm, X.14
 LSB, II.31, III.12
 in comparator, IV.3
 size, fullscale converter, XI.3

M

MAC, VIII.9-10
 DMD bus, VIII.9
 features, VIII.10
 instructions, VIII.11
 PMD bus, VIII.9
 register, VIII.9
 Magnetic field, shielding, XI.28-29
 Mahoney, Matthew, III.44
 Mar, Amy, VII.32
 MASH:
 cascaded stable first-order loops, VI.26
 digital differentiator, VI.27
 modulator, VI.27
 sigma-delta ADC, VI.26
 topology, VI.27
 Matsuya, Y., VI.30
 Maximally flat filter, II.39
 MDAC, V.15
 4-quadrant, V.16
 configuration, V.16
 12-bit, features, V.17
 16-bit linear,
 attenuation, V.17
 selected inpput codes, V.17
 applications, V.18
 bipolar input, V.15
 construction, V.16
 Mecca ground, XI.36
 MECL System Design Handbook, XI.63
 Meehan, Pat, III.43
 Memory read, See: RD
 Metal lids, on ceramic IC package, XI.18
 Microcomputer, ADSP-2101, VIII.5
 Microstrip:
 controlled impedance, mixed signal system,
 XI.44
 transmission line, XI.41
 Minimum 4-Term Blackman-Harris window,
 VII.29, 31
 Mixed signal circuit, XI.1
 design, XI.2
 prototyping, XI.61
 Mixed signal processing, I.1-2, X.1
 definition, I.5
 Mixed signal system:
 ground plane, XI.44
 signal routing, XI.44
 Mode status register, See: MSTAT
 Modem, V.1, X.1-3
 CCITT recommendations, X.3
 demodulation, X.1
 echo cancellation, X.1
 error detection and correction, X.1
 full-duplex, X.2
 half-duplex, X.2
 hybrid circuit, X.2
 modulation, X.1
 modulation methods, X.3
 V.32,
 adaptive equalizer, X.7-8
 ADC, X.5, 7-10
 analog front end, X.9-10
 antialiasing filter, X.8, 10
 characteristics, X.4
 CODEC, X.10
 DAC, X.5-6, 9
 demodulator, X.7-8
 descrambler, X.7-8

diagram, X.5
differential decoder, X.7-8
echo canceller, X.7-8
error correction, X.9
filter chips, X.10
FIR filter, X.9
GSTN, X.4
I/O port, X.8, 10
input antialiasing receiver, X.7-8
modulation, X.5-6
pulse shaping filter, X.7
QAM, X.5
receiver, X.7-8
SCF, X.10
signal constellation, X.6
signal separation, X.8
SNR, X.10
THD, X.10
transmitter, X.5-6
Viterbi decoder, X.7-9
Modulation:
 ASK, X.3
 FSK, X.3
 PSK, X.3
 QAM, X.3
Modulator, II.15-6
 as precision rectifier, II.16
Momentum Data Systems, Inc., VII.32
Morley, Nick, X.30
MSB, III.33
MSTAT, VIII.14
Multi-bit sigma-delta converters, VI.27
Multi-stage noise shaping, See: MASH
Multiple feedback, filter configuration, II.40-43
Multiplier:
 2-quadrant linear, II.11
 4-quadrant, II.15
 translinear, II.12
 as audio power meter, II.13
 basic, II.8
 basic transconductance, II.10
 direct divide capable, II.15
 as divider, II.14

 Gilbert cell, II.11
 in op-amp, II.14
 as power meter, II.13
 signal output, II.15
 trimmable error, II.12
 types, II.9
 VGA, II.17
Multiplier-accumulator, See: MAC
Multiplying DAC, See: MDAC
Murphy's Law, XI.1
 corollary, XI.2

N

N-Bit flash converter, IV.3
Nahman, Norris S., III.43
Net midscale glitch, III.33-34
Non-inverting adder, II.2
NRZ data, II.5, 7
Nyquist bandwidth, III.8-9, 12-13, 23-24
Nyquist criterion, III.2-3, 7, 9
Nyquist limit, III.35, 42
Nyquist sampling, analog lowpass filter, III.8, VI.2
Nyquist theorem, III.37

O

Ohm's Law, XI.2
Ohm's law, PCB copper foil, XI.4
Op-amp, I.7, III.37, IV.4, V.4
 absolute value circuit, II.4
 adder, II.1
 basic configuration, I.9
 configuration, I.8-9
 differentiator, II.2
 dynamics, II.48, IV.4
 external, for buffered output voltage, II.31
 filters, II.37
 high-performance, II.1
 ideal equations, I.9
 instrumentation, I.10-11
 integrator, II.3
 inverting, I.9
 level shifter, I.9

- ideal equations, I.9
- non-inverting, I.9
- performance, III.22
- specifications, I.8
- subtractor, II.2
- Operational amplifier, See: Op-amp
- Oppenheim, A.V., VII.32
- Oscillator:
 - Colpitts, XI.57
 - noise, XI.57
- Oversampling:
 - analog filter, III.9, VI.3
 - digital filter, III.9, VI.3
 - ratio, III.8-9
- Overvoltage recovery time, III.30

P

- PAL color subcarrier frequency, I.15
- Pan-European Digital Cellular Radio System (Groupe Speciale Mobile), See: GSM
- Parallel DAC, DSP interface, IX.10
- Parallel peripheral device, read interface, IX.3
- Parks-McClellan algorithm, VII.12
- Parzefall, F., VI.30
- Passive RC network, SC equivalent, III.39
- PC track, transmission line, XI.40
- PCB:
 - capacitance, XI.15-16
 - differential transmission, XI.46
 - edge connections, XI.45
 - flowchart layout, XI.45
 - ground error minimization, XI.46
 - as ground plane, XI.39
 - grounding, XI.42
 - guard ring, XI.7
 - multiple card system,
 - ground, XI.43
 - star analog ground, XI.43
 - multiple ground pins, XI.46
 - prototyping, XI.61-62
 - SPICE modeling, XI.59
 - Teflon stand-off insulator, XI.7-8
 - thermoelectric effects, XI.12
 - transmission line, XI.40
- PCB copper foil, XI.3
 - anisotropy, XI.6
 - insulation, XI.6
 - Ohm's law, XI.4
 - skin effect, XI.4-5
 - track resistance, XI.4
- Peak detector, II.4-6
- Peak harmonic content, III.23
- Peak spurious component, III.23
- Peak spurious response, III.23
- Phase shift keying, See: PSK
- Phase-locked loop, in disk drive read amplifier, II.6-7
- Phased array, digital, II.34
- Phased array ultrasound, II.34
- Physics, laws, XI.2
- PMA bus, ADSP-2101 microcomputer, VIII.7
- PMD bus:
 - ADSP-2101 microcomputer, VIII.7
 - MAC, VIII.9
- PMS, DSP parallel interface, IX.1
- Pohlmann, Ken, X.30
- Polarity, in analog multiplier, II.8
- Pole, filter, II.40
- Potentiometer, V.15
- Power supply:
 - for electronic circuit, XI.47
 - noise, XI.47-48
 - switching-mode, XI.48-49
 - noise, XI.48
- Precision analog circuit, ESD, XI.9
- Precision rectifier, II.15-16
- Pressure measurement, sigma-delta ADC, VI.16, 19
- Printed circuit board, See: PCB
- Process control, I.3
- Processor interrupt request, See: IRQ
- Program memory address, See: PMA
- Program memory data, See: PMD
- Program memory select, See: PMS
- Program sequencer, VIII.14

ADSP-2101 microcomputer, VIII.15

ASTAT, VIII.15

features, VIII.15

ICNTL, VIII.15

IFC, VIII.15

IMASK, VIII.15

MSTAT, VIII.15

SSTAT, VIII.15

Programmable gain amplifier, V.15, 18

Programmable Lowpass Analog Filter Using
12-Bit DACs, II.51

Programmable power supply, V.18

Programmable state variable filter, II.48

Prototyping, XI.61-62

PSK, modulation, X.3

Public Switched Telephone Network, See:
GSTN

Q

Q:

inductor, XI.32

resonant circuit, XI.32

QAM, DAC, X.7

Quadrature amplitude modulation, See:

QAM

Quality factor, See: Q

Quantization, III.12-13

Quartz crystal, Q, XI.32

R

Rabiner, Lawrence R., III.44, VII.32

Radix-2 decimation-in-time, FFT butterfly,
VII.26

Radix-2 FFT, VII.24-26

benchmark processing time, VII.28

Radix-4 FFT, VII.24-25

algorithm, VII.26

benchmark processing time, VII.28

RAM, fixed boundary, VIII.4

RAMDAC, I.4

Ramirez, Robert W., III.43, VII.32

RC filter:

SCF equivalent, III.37, 39

single-pole passive, III.37, 39

RD:

ADSP-2101, IX.4

DSP, IX.1

RD timing, ADSP-2101, IX.2

Real-world signal, processing, I.3

Receive frame synchronization, See: RFS

Receiver, balanced, I.13-15

Rectifier, full wave, II.4

Recursive subranging, IV.7

Reduced-instruction-set computer, See: RISC

Reidy, John, III.43, X.30

Remez exchange algorithm, VII.12

Resistance, conductor, XI.3

Resistive system, logic fan-out, XI.53-54

Resistor:

equivalent circuit, XI.11

inductance, XI.11

Johnson noise, XI.14-15

mismatch, XI.13

parasitic effects, XI.11

induction, XI.11

thermal noise, XI.14

thermal stability, XI.13

thermoelectric effects, XI.12

voltage variation, XI.13-14

wirewound, thermocouple, XI.12

Resonant circuits, from decoupled power
lines, XI.30

RFS, DSP, IX.15

Ribbon cable:

mutual inductance, XI.28

signal coupling, XI.28

Ringling, XI.29

RISC:

application, VIII.1

architecture, VIII.1-2

DSP, VIII.1

programming, VIII.1

RMS, II.18

RMS conversion:

explicit, II.19

implicit, II.20

wideband, II.19

RMS quantization, III.13
 noise, III.13
 RMS value:
 explicit method, II.18-19
 implicit method, II.18, 20
 Root mean square, See: RMS
 Rosso, M., X.30

S
 S/N+D, III.21-22
 Sallaerts, D., X.30
 Sallen-Key filter, II.40, 43, 50
 frequency response, II.50
 Sample-and-hold amplifier, See: SHA
 Sampled data system, III.1
 Sampling clock, design, noise minimization, XI.56
 SAR, IV.1
 ADC encoder, IV.3
 ADC with SHA, IV.2
 clock-to-data-output delay, IV.3
 input data setup time, IV.3
 SAW, III.37
 SCF, III.37, 40
 bandwidth, III.38
 sampling device, III.37
 Schafer, R.W., VII.32
 Schoenwetter, Howard K., III.43
 Serial port:
 features, VIII.17
 interface lines, VIII.16
 Settling time, III.29, 31-32
 Seven-pole FDNR 20kHz antialiasing filter, II.49
 SFDR, III.23
 SHA, II.5-6, III.21-22, 24, IV.1, 3-4, 7
 degitcher, III.35-36, V.13-15
 switching transients, V.13
 SHA-ADC pair, III.21
 Shaped quantization noise, and digital filtering, VI.9
 Sheingold, Daniel H., I.18, II.51, III.43, V.19
 Shifter, VIII.11

 features, VIII.12
 instructions, VIII.12
 operations, VIII.12
 SID, in receiver, X.16
 Sigma-delta ADC, VI.1
 16-bit, VI.16-17
 calibration, VI.17
 digital filter response, VI.18
 digital filter step response, VI.18
 mixed signal device, VI.24-25
 modulator, VI.17
 18-bit audio, VI.13
 amplitude response, VI.15
 antialiasing filter, VI.13
 digital filter, VI.15
 modulator output, VI.14
 specifications, VI.14
 21-bit,
 digital filter response, VI.20
 features, VI.19-20
 high-level analog input, VI.21
 on-chip signal conditioning, VI.19
 single-channel differential low-level PGA input, VI.21
 two-channel differential low-level PGA input, VI.21
 22-bit,
 configuration, VI.19
 low frequency measurement, VI.19
 DAC, VI.3, 27
 decimation, VI.1, 8-10
 digital filtering, VI.1, 8-10, 16
 filter output, VI.8
 FIR filter, VI.8, 10
 first-order, VI.4
 noise shaping, VI.6-7
 first-order modulator,
 idling pattern, VI.11
 integrator output, VI.11
 higher order modulator loops, VI.12
 idling pattern, VI.11
 IIR filter, VI.8, 10
 low frequency measurements, VI.16, 19
 MASH, VI.26

- for mixed signal VLSI, VI.1
- modulator, VI.3-4
 - bit pattern, VI.11
 - frequency domain model, VI.5
 - second-order, VI.6-7
 - waveforms, VI.4
- multi-bit, first-order, VI.27
- noise-spectrum shaping, VI.1-2
- Nyquist sampling, VI.2
- oversampling, VI.1-3, 8, VII.19
- quantization noise shaping, VI.3
- second-order, VI.6-7
 - noise shaping, VI.6-7
- second-order modulator,
 - idling pattern, VI.12
 - second integrator output, VI.12
- shaped quantization noise, VI.9
- shaped quantization noise distribution, VI.6
- SNR, VI.6, 8
- synchronous voltage-to-frequency converter, VI.5
 - tonal considerations, VI.11
 - two-stage filter, VI.10
- Sigma-delta architecture, IX.27
- Sigma-delta CODEC:
 - decoder, VI.25-26
 - linear coded ADC and DAC, VI.25
 - mixed signal device, VI.24-25
- Sigma-delta converter, III.7
 - limits, VI.28
 - MASH, VI.26
 - multi-bit, VI.27
 - multi-bit versus single bit, VI.28
 - oversampling, III.7
 - summary, VI.29
- Sigma-delta DAC:
 - 16-bit,
 - analog filter, VI.23
 - digital modulator, VI.23
 - discrete time signal, VI.24
 - IIR filter, VI.23
 - mixed-signal device, VI.24-25
 - shaped quantization noise, VI.23
 - voiceband output, VI.23
 - antialiasing filter, VI.21
 - in compact disk players, VI.22
 - concepts, VI.22
 - deglitcher, VI.23
 - glitches, VI.22
 - lowpass filter, VI.22
 - mixed-signal IC architecture, VI.21
 - noise filtering, VI.21
 - oversampling, VI.21
- Signal:
 - ADC, I.2
 - analog, I.1
 - attenuation, II.32
 - autocorrelation, I.3
 - characteristics, I.1
 - conditioning, I.1, 8
 - op-amp, I.7
 - convolution, I.3
 - definition, I.1
 - digital, I.1
 - dynamic range, for ADC selection, III.16
 - filtering, I.3
 - magnitude, II.18
 - origin, I.1
 - output, of multiplier, II.15
 - processing, I.2
 - analog, See: Analog signal processing
 - analog vs. digital, I.5-7
 - frequency content compression, I.2-3
 - real-time, I.5
 - real-world, I.2
 - ASP, I.4
 - digitally synthesized, I.4
 - DSP, I.4
 - generation, I.4
 - MSP, I.4
 - processing, I.3-4
 - TDM, I.2
- Signal lead, voltage drop, XI.6
- Signal routing, XI.32, 44
 - mutual inductance, XI.26-27
- Signal voltage:
 - Mecca ground, XI.36

- ground, XI.36
 - signal-to-noise plus distortion, See: S/N+D
 - signal-to-noise ratio, See: SNR
 - signal descriptor, See: SID
 - silicon junction transistor, transconductance, 10
 - sinewave curvefit, for ENOBs, III.15
 - $\sin(x)/x$ frequency rolloff effect, III.35-36
 - Skin effect:
 - DC circuit, XI.4
 - PC board dielectric, XI.5
 - PCB copper foil, XI.4-5
 - Slattery, Bill, XI.63
 - Slope voltage, II.21
 - Sluyter, R.J., X.30
 - SNR, III.12, 21, 24, 27, 33
 - calculation, III.13-15
 - ENOBs, III.13
 - measurement, III.21-22
 - Socket:
 - capacitance, XI.60
 - high performance analog circuit, XI.60
 - inductance, XI.60
 - resistance, XI.60
 - undesirable in circuit, XI.60
 - Sodini, C.G., VI.30
 - Soft-body tissue attenuation, II.32
 - Southcott, C.B., X.30
 - Spectral analysis, VII.23
 - Speech CODEC, X.14
 - SPICE, modeling, limitation, XI.58-59
 - Spurious free dynamic range, See: SFDR
 - SSTAT, VIII.14
 - Stack status register, See: SSTAT
 - Star ground, XI.36
 - State variable, filter configuration, II.40-43
 - Stray capacitance, chip bondwires, XI.19
 - Subtractor, II.3
 - instrumentation amp, II.2
 - op-amp, II.2
 - Successive approximation architecture, IX.27
 - Successive detection log amp, II.21
 - Summing amplifier, II.26
 - Super-Nyquist, III.9, 11
 - sampling, FDM signal, III.10
 - Swanson, E.J., VI.30
 - Switched capacitor:
 - CMOS, III.39
 - resistor, III.38
 - Switched capacitor filter, III.37
 - advantages, III.40
 - limitation and error source, III.42
 - as resistor, III.37, 39-40
 - Switched capacitor integrator, III.41
 - System ground, XI.42
 - System interface, ADSP-2101 microcomputer, VIII.17
- ## T
- T-Carrier, digital voice transmission, X.14
 - Tant, M.J., III.44
 - TDM signal, I.2
 - TDMA, DSP, signal transmission, X.13
 - Telephone:
 - diagram, X.2
 - imperfections, X.1
 - long distance transmission, X.2
 - Television transmultiplexer, III.9
 - Temperature measurement, sigma-delta ADC, VI.16, 19
 - TGA, II.32-34, 36
 - THD+N, for 20-bit audio DAC, V.12
 - THD, III.21, 23, 33-35
 - Thermometer code, IV.4
 - Time division multiple access, See: TDMA
 - Time division multiplexed, See: TDM
 - Time gain amplifier, See: TGA
 - TMS32020/C25, serial interface, IX.19, 22
 - Total harmonic distortion, See: THD
 - Track-and-hold, for deglitching, V.13
 - Transcoder, X.14
 - Transducer, slow, IV.8
 - Transient response, III.29
 - Transimpedance amplifier, II.15
 - Translinear multiplier, II.11, 13, 15-16
 - Transmission line, XI.40
 - Transmit data, See: DT

True log amp, II.21, 31

TTL:

logic fan-out, XI.53-54

noise, XI.57-58

Tukey, J.W., III.43

Tuned circuit, XI.29

from stray capacitance, XI.31

U

Ultrasound, II.32-33

amplifier, II.34-35

phased array, II.34

VGA, II.36

Undersampling, III.9

V

VAD, DTX, X.16

Van Valkenburg, M.E., II.51

Variable gain amplifier, See: VGA

Vary, P., X.30

Very large scale integration, See: VLSI

VGA, II.5-6, 15, 17-18, 32-33, 35-6

AC response, II.18

independent gating, II.36

multiplier, II.17

single channel, II.36

transient response, II.18

Video amp, II.21, 23

Video line driver, low differential gain and phase, I.16

Video signal transmission, transmission, XI.47

Viterbi decoding, X.9

VLSI, I.5, IV.3, 9

Voice activity detection, See: VAD

Voice encoding, CODEC, X.14

Voiceband filter, III.38

Voltage:

RMS value, II.18

See also EMF

Voltage drop:

resistive, gain error, XI.6

sense connection, XI.6

Voltage-controlled attenuator, II.36

W

Wainwright Instruments GmbH, XI.63

Wainwright Instruments Inc., XI.63

Weeks, Pat, III.43

Weigh scales, sigma-delta ADC, VI.16, 19

Weighting function, VII.31

Welland, D.R., VI.30

Wideband amplifier, I.12

variable gain low noise, II.35

Wideband RMS measurement, II.19

Wideband Sallen-Key filter, II.50

Williams, A.B., II.51

Williams, C.S., VII.32

Wooley, Bruce, VI.30

WR, DSP, IX.8

Write output, See: WR

Wynne, John, XI.63

Z

Zeoli, G.W., III.43

Zverev, A.I., II.51

MIXED SIGNAL PROCESSING DESIGN SEMINAR

ADSP-2111, VI.25
ADSP-21msp50, X.14, 16-18
ADSP-28msp01, X.10-11
ADSP-28msp02, VI.24-25, IX.1, 23, 25-6,
X.10

S

SSM-2141, 1.14-15
SSM-2142, 1.14-15

INDEX

V

AD1865, X.25-27
AD1876, X.23
AD1879, VI.1, 12-15, 26-28, X.23
AD7111, V.17-18
AD7341, X.9-10
AD7371, X.9-10
AD7528, II.48, 51
AD7701, VI.16-18
AD7703, VI.17, 19
AD7710, VI.19-21
AD7711, VI.19-21
AD7712, VI.19-21
76-8, IX.10-13

INDEX-24



WORLDWIDE HEADQUARTERS

One Technology Way, P.O. Box 9106, Norwood, MA 02062-9106, U.S.A.
Tel: (617) 329-4700, Fax: (617) 326-8703, Telex: 924491